arXiv:1410.8766v1 [stat.ME] 31 Oct 2014

Local Neighborhood Fusion in Locally Constant Gaussian Graphical Models

Apratim Ganguly and Wolfgang Polonik

Apratim Ganguly
Postdoctoral Associate
Department of Mathematics and Statistics
Boston University
111 Cummington Mall
Boston, MA 02215
e-mail: apratimganguly@gmail.com
url: http://math.bu.edu/people/apratim

Wolfgang Polonik
Professor
Department of Statistics
University of California, Davis
One Shields Ave.
Davis, CA 95616
e-mail: wpolonik@ucdavis.edu
url: http://anson.ucdavis.edu/ polonik

Abstract: In this paper we penetrate and extend the notion of local constancy in graphical models that has been introduced by Honorio et al. (2009). We propose *Neighborhood-Fused Lasso*, a method for model selection in high-dimensional graphical models, leveraging locality information. Our approach is based on an extension of the idea of node-wise regression (Meinshausen-Bühlmann, 2006) by adding a fusion penalty. We propose a fast numerical algorithm for our approach, and provide theoretical and numerical evidence for the fact that our methodology outperforms related approaches that are ignoring the locality information. We further investigate the compatibility issues in our proposed methodology and derive bound for the quadratic prediction error and l_1-bounds on the estimated coefficients.

Primary 62J07, 60G60; secondary 62H20, 62F12.
Keywords and phrases: Graphical Models, Gaussian Graphical Models, Neighborhood Selection, Model Selection, Local Constancy, Lasso, Fused Lasso.

1. Introduction

In the context of a graphical model, the concept of local constancy has been coined recently in Honorio et al., 2009. The goal of the present paper is to penetrate this idea by providing a thorough analysis and an extension of this concept. Our approach is using the idea propounded by Meinshausen and Bühlmann (2006) for structure learning in Gaussian graphical models, and it incorporates an additional fusion penalty term that aims to enforce the structural constraint

*Funding for this work was provided by NSF grant DMS-1107206.

of *local constancy*. When the nodes of a graph have spatial information attached, then local constancy can be interpreted as a certain type of spatial regularity (see below). Conceptually, the entire neighborhood graph given by the graphical model is split into two subgraphs: (a) A *local* graph consisting of nodes that are spatially close, and (b) a *non-local* graph, where the edges connect nodes of two different spatial clusters of nodes. Our approach then assumes that we have prior structural knowledge about the local graph that we aim to incorporate into the estimation/model selection approach.

It is well known that a p-dimensional graphical model is given by a p-dimensional Gaussian distribution with non-singular covariance matrix Σ of a random vextor X. The conditional independence of the components of X can be represented by occurrence of zero entries of precision matrix $\Sigma^{-1} = \Omega$. The non-zero entries of the precision matrix define the edges of the graph corresponding to this multivariate Gaussian distribution.

Neighborhood selection algorithms aim to find all the neighbors of a node X_a in a graphical model based on an i.i.d. sample. Meinshausen & Bühlmann [18] showed that this problem can be interpreted as an ensemble of l_1-penalized regressions, each of which can be solved using the lasso algorithm of Tibshirani (1997). Our proposal is an extension of this approach. In order to incorporate this additional structural constraint we assume that the prior structural knowledge can be translated into a *local* neighborhood graph, which we think of being comprised of regular graph objects such as chains, cycles, lattices or cliques. The availability of additional locality information is critical when data is observed in a certain manifold with spatial geometry. The assumption of local constancy, in some sense, enforces spatial regularization on structure learning and thereby stimulates the search for probabilistic dependencies between local clusters of nodes. We would like to emphasize that the knowledge of this prior structural information is based on domain knowledge and hence known beforehand.

We propose to use the Meinshausen & Bühlmann approach and to add a new penalty term which, in essence, generalizes the fused lasso penalty (see Tibshirani, 2005) by extending the prenalty term over differences of nodes in the local neighborhood graph, which in turn is given by the prior structural information. This leads to the *neighborhood-fused lasso* procedure, a model selection method for locally constant Gaussian graphical models. We provide theoretical and numerical evidence that our approach outperforms competing model selection algorithms where locality information is ignored.

The paper is organized as follows: In section 2 we first provide an overview of related work. Sections 3 and 4 then provide a precise definition of our approach descibed above, and a discussion of local constancy, respectively. In section 5 we propose our optimization algorithm, before we present both both finite sample and large sample properties of our appraoch in section 6. We prove theoretically that introducing a local penalty term reduces the finite sample type-I error probability in model selection and leads to equivalent accuracy with smaller sample size than competitors. We also discuss data dependent choice of the penalty parameters, which avoids the use of cross-validation based methods. We study the asymptotic l_1 properties of our estimator to find sufficient conditions

on design matrix and regularization parameters to find nice asymptotic bounds in parameter estimation and prediction in terms of l_1 and l_2 metric, respectively. In particular, our theoretical analysis reveals that with our assumptions our proposed method displays all the desired properties like sign-consistency and model selection consistency. Some numerical results can be found in section 7. The proofs are delegated to the appendix.

2. A Review of Related Work

Going back to Dempster [8] who introduced *Covariance Selection* to discover the conditional independence restrictions (the graph) from a set of i.i.d. observations, many methods have been proposed for sparse estimation of the precision matrix in a Gaussian graphical model. The proposed procedures usually rely on optimization of an objective function [9, 16]. While modern technological developments and high computing power enable us to deal with high dimensional models, there still are computational challenges. Usually, greedy forward-selection or backward-deletion search is used. In forward (backward) search, one starts with the empty (full) set and adds (deletes) edges iteratively until a suitable stopping criterion is fulfilled. The selection (deletion) of an edge requires an MLE fit [21] for $O(p^2)$ many models, making it a suboptimal choice for high-dimensional models where p is large. Also, the MLE might not exist in general for $p > n$ (see [4]). In contrast, neighborhood selection using lasso, as proposed by Meinshausen and Bühlmann [18], relies on optimizing a convex function applied consecutively to each node in the graph, thus fitting $O(p)$ many models. Fast lasso-type algorithms and data dependent choices for regularization parameter reduce the computational cost. Unlike covariance selection, this algorithm estimates the dependency graph by sequential estimation of individual neighbors and subsequent combination by taking unions or intersections. Other authors have proposed algorithms for the exact maximization of the l_1-penalized log-likelihood. Yuan & Lin [28] proposed an l_1-penalty on the off-diagonal elements of the concentration matrix for its sparse estimation with the positive definiteness constraint. They showed that this problem is similar to the *maxdet* problem (see Vandenberghe et al. [26]), and thus solvable by the interior point algorithm. A quadratic approximation to the objective function in their proposed method leads to a solution similar to Meinshausen & Bühlmann. Banerjee et al. [1] viewed this as a penalized maximum likelihood estimation problem with the same l_1-penalty on the concentration matrix. Constructing the dual transforms the problem into sparse estimation of the covariance matrix instead of the concentration matrix. They proposed *block coordinate descent algorithm* to solve this efficiently for large values of p. They also showed that the dual of the quadratic objective function in the block-coordinate step can be interpreted as a recursive l_1-penalized least square solution. Friedman, Hastie & Tibshirani [10] used this idea successfully to develop an algorithm known as *graphical lasso*. Using a coordinate descent approach to solve the lasso problem speeds up the algorithm to a considerable extent, making it quite fast and effective for a large class of high-dimensional problems.

In all the aforementioned methods, information on the local geometry is not taken into consideration. Often one encounters data that are measured on a certain manifold. E.g., data describing some feature on the outline of a (moving) silhouette, or pixels or voxels in an 2-d or 3-d image, respectively. In most of these problems, spatially close variables have a structural resemblance in terms of probabilistic dependence. Exploiting this local behavior might lead to faster and/or more efficient estimation. Honorio, Ortiz, Samaras et al. [14] introduced the notion of *local constancy*. According to them, if one variable X_a is conditionally (in)dependent of X_b, then a local neighbor of X_a in that manifold is likely to be conditionally (in)dependent of X_b as well. They developed a*coordinate direction descent algorithm* to solve a penalized MLE problem that penalizes both the l_1-norm of the precision matrix and the l_1-norm of local differences of the precision matrix, expressed as its "diagonal excluded product" with a local difference matrix. Local geometry has been addressed, although implicitly, by Tibshirani et al. [23] in the context of fused lasso. Chen et al. [6] used this idea and proposed *graph guided fused lasso* for structure learning in multi-task regression, where the output space is continuous and the outputs are related by a graph. The input space is high dimensional and outputs that are connected by an edge in the graph are believed to share a common set of inputs. The goal then is to learn the underlying functional map from the input space to the output space in a way that respects the similar sparsity pattern among the covariates that are believed to affect the output variables that are connected. Local smoothing by penalizing differences of neighboring nodes has been discussed in Kovac and Smith [15] in the context of nonparametric regression based on observations on nodes of a graph. Their algorithm aims to split the image into active sets and subsequently merge, split or amalgamate them in order to minimize a penalized weighted distance.

3. Neighborhood selection using Fused Lasso

Like Greenshtein & Ritov [11] and Meinshausen & Bühlmann [18] we work in a set-up where the number of nodes in the graph $p(n)$ and the covariance matrix $\Sigma(n)$ depend on the sample size. Consider the $p(n)$-dimensional multivariate random variable $X = (X_1, \cdots, X_p) \sim N(\mu, \Sigma)$. The conditional (in)dependence structure of this distribution can be represented by the graph $\mathcal{G} = (\Gamma(n), E(n))$, where $\Gamma(n) = \{1, \cdots, p(n)\}$ is the set of nodes corresponding to each coordinate variable and $E(n)$ the set of edges in $\Gamma(n) \times \Gamma(n)$. A pair of nodes (a, b) lies in $E(n)$ if and only if X_a is conditionally dependent of X_b, given all other remaining variables $X_{\Gamma(n)\setminus\{a,b\}} = \{X_k : k \in \Gamma(n)\setminus\{a, b\}\}$.

The neighborhood ne_a of a node a is defined as the smallest subset of $\Gamma(n) \setminus \{a\}$ such that given ne_a, X_a is conditionally independent of all the remaining nodes. In other words, neighbors of a certain node consists of the coordinates that are conditionally dependent on that particular node. As already mentioned, the conditional independence can be represented by occurrences of zero entries at respective cells of precision matrix Ω, i.e., $X_a \perp X_b | X_{\Gamma(n)\setminus\{a,b\}}$ iff $\Omega_{ab} =$

$\Sigma^{-1}_{ab} = 0$. The neighborhood selection / model selection algorithms aim to find all the neighbors of a node X_a, given n i.i.d. observations of X. Meinshausen & Bühlmann considered this as a penalized regression problem, where each variable is regressed on the remaining variables with an l_1 penalty on the estimated coefficients.

By the above definition of neighborhood, we have, for all $a \in \Gamma(n)$ that

$$X_a \quad \perp \quad \left\{X_k : k \in \Gamma(n) \setminus \{a \cup \mathrm{ne}(a)\}\right\} \ \mid \ X_{\mathrm{ne}(a)}.$$

An alternative definition of neighborhood of X_a is given as the non-zero components of θ^a where θ^a is given by

$$\theta^a \ = \ \mathrm{argmin}_{\theta:\theta_a=0}\mathbb{E}\left(X_a - \sum_{k\in\Gamma}\theta_k X_k\right)^2. \tag{3.1}$$

In light of the above definition, the set of neighbors of a node $a \in \Gamma(n)$ is

$$\mathrm{ne}(a) = \{b \in \Gamma(n) : \theta^a_b \neq 0\}. \tag{3.2}$$

Given the domain knowledge we construct a regular graph $\mathcal{G}_{\mathrm{local}}$ that is representative of the underlying spatial geometry. We call it a *local neighborhood graph*. For example, $\mathcal{G}_{\mathrm{local}}$ could be a chain graph, a two or three dimensional lattice or more generally a collection of cliques or wheels. Let E_{local} denote the edge set corresponding to $\mathcal{G}_{\mathrm{local}}$, then define:

Definition 3.1. X_a and X_b are *local neighbors* with respect to $\mathcal{G}_{\mathrm{local}}$ if the edge connecting them belongs to E_{local}.

Note that the edge between X_a and X_b need not belong to $E(n)$. Now if one incorporates the *local constancy* property to this graph, the neighborhood becomes more structured. By local constancy, if X_b is conditionally independent of X_a given the other nodes (and hence, there is no edge between a and b), it is likely that a *local neighbor* $X_{b'}$ of X_b is also conditionally independent of X_a, making both θ^a_b and $\theta^a_{b'}$ equal to zero. Thus, the zeroes of θ^a are expected to reflect the sparsity pattern.

We now generalize the traditional fused lasso algorithm [23] to satisfy our purpose of estimating the zeros and non-zeros of the precision matrix by exploiting the locality information given by the local neighborhood graph $\mathcal{G}_{\mathrm{local}}$. Before doing that we carefully define the *difference matrix*. Consider an $m \times p$ matrix D where $m = |E_{\mathrm{local}}|$ is the total number of pairs of local neighbors. Assuming that we have a labelled sequence of nodes $\Gamma(n) = \{1, 2, \cdots, p(n)\}$, arrange the pairs of local neighbors in a sequence

$$\mathcal{B} := \{\{(u, v_u) : v_u \in \mathrm{ne}_l(X_u) > u\} : u \in \Gamma(n)\},$$

where $\mathrm{ne}_l(X_u)$ denotes the set of local neighbors of X_u. The inequality $v_u > u$ is included to ignore double counting. Note that $\mathcal{B}$ is nothing but a convenient ordering of edges in $\mathcal{G}_{\mathrm{local}}$, and hence $\mathcal{B}$ contains the same information as $\mathcal{G}_{\mathrm{local}}$.

It should also be mentioned that our results are not influenced by the particular labelling used. The k^{th} row is given by $D_{k,.} = e_i - e_j$ where (i, j) is the k^{th} element of the local neighbor sequence and e_i, e_j denote two canonical basis vectors of $\mathbb{R}^p$ where the 1's occur at i^{th} and j^{th} position, respectively. This way each pair is represented by a row in the difference matrix. We denote by D^a the $m_a \times p$ sub-matrix of D selecting all the rows with a^{th} entry being 0. In other words, D^a is the difference matrix corresponding to all the local neighbor pairs not involving X_a. The number of local neighbors of X_a in $\Gamma(n) \setminus \{a\}$ is m_a. It should be noted that throughout our discussion we shall assume that the local neighborhood structure is known to us, meaning that D is known beforehand and it does not depend on the data.

Now we define the neighborhood-fused LASSO estimate $\hat{\theta}^{a,\lambda,\mu}$ of θ^a.

$$\hat{\theta}^{a,\lambda,\mu} \quad = \quad \text{argmin}_{\theta:\theta_a=0} \left(n^{-1} ||X_a - X\theta||^2 + \lambda ||\theta||_1 + \mu ||D^a \theta||_1 \right), \tag{3.3}$$

where $||x||_1$ denotes the l_1-norm of x. Penalizing both the l_1-norms of θ and $D^a\theta$ implies parsimony, thus ensuring sparsity and local constancy at the same time. This property helps us in variable selection and thereby leads to neighborhood selection. Note that we are referring to both the non-local and the local neighbors. The neighborhood estimate of node a is defined by the nodes corresponding to non-null coefficients when X_a is regressed on the rest of the variables with the fused lasso penalty in (3.3), in other words

$$\hat{\text{ne}}_a^{\lambda,\mu} \quad = \quad \left\{ b \in \Gamma(n) \setminus \{a\} : \hat{\theta}_b^{a,\lambda,\mu} \neq 0 \right\}, \tag{3.4}$$

where $\hat{\theta}_b^{a,\lambda,\mu}$ denotes the b-th component of the vector $\hat{\theta}^{a,\lambda,\mu}$. It is clear that the selected neighborhood depends on the value of λ and μ chosen. Large values of λ and μ will give rise to more sparse solutions. Usually the regularization parameters are chosen by some cross validation criteria. But in this paper we will find a data driven approach for selection of regularization parameters that speeds up computation and ensures asymptotic consistency in model selection. Meinshausen and Bühlmann derived a data driven choice for λ in their paper. We will extend their method for simultaneous selection of λ and μ from the data.

4. A Discussion of Local Constancy

Here we discuss further formalizations of the notion of local constancy. It should be noted again that the definition is linked with the ideas used by Honorio et al. [14], but our definition is more general in nature. First we introduce the notion of *diagonal excluded matrix product*. Given a matrix A, its zero operator is defined as $Z(A) := (\mathbb{I}(A_{ij} = 0))_{i,j}$ where $\mathbb{I}$ is the indicator function. The diagonal excluded product of two matrices A and B can now be defined as,

$$A \oslash B := Z(A) \circ AB,$$

where $\circ$ denotes the Hadamard product of matrices. Although the name does not clearly show how diagonals are removed from the product, usually if A is taken to be a difference matrix defined above, this will eventually lead to the exclusion of diagonals of the matrix B.

4.1. A more quantitative perspective

First observe that the absolute values of the non-zero entries of $D \oslash \Omega$ are of the form $|\omega_{jk} - \omega_{jk'}|$ where (k, k') corresponds to a pair of local neighbors. In fact, we have $\|D \oslash \Omega\|_1$ equals the sum of all the absolute values of the differences $\omega_{ik} - \omega_{jk}$, $k \notin \{i, j\}$ for local neighborhood pairs (i, j). In other words, if we think of the entires ω_{ik} as measuring conditional dependence of two nodes, then $\|D \oslash \Omega\|_1$ measures in how different the local neighborhood pairs are with respect to conditional dependence. A small value of $\|D \oslash \Omega\|_1$ indicates that the local neighbors behave similarly in this respect.

Thus, one way to define the local constancy property is to impose a bound on the norm of the difference of local neighbors. We say a model exhibit $(\epsilon, \mathcal{G}_{\text{local}}, l_p)$-local constancy if

$$\left\|D \oslash \Omega\right\|_p < \epsilon$$

where $\left\| \cdot \right\|_p$ denotes the l_p norm. For our purpose, the best candidate would be $p = 1$ for certain desirable properties like sparsity and convexity. However, a different choice of norm would lead to a different solution which might be more appropriate for other types of problems. With $p = 1$, this definition coincides with Honorio's local penalty criterion.

It should be noted that the ϵ parameter controls the degree of local constancy so this is closely related to the tuning parameter we are going to use. This serves as a constraint on certain matrix parameters that we are trying to estimate. A small ϵ would ensure a high level of local constancy where a large ϵ would do the opposite.

4.2. A Bayesian Perspective

Since the locality information is being used here as a prior information, it makes sense to interpret it from a Bayesian perspective. Let $\pi(\Omega)$ denote a probability distribution on the space of positive definite matrices. Then one can say that an $(\epsilon, \delta, \mathcal{G}_{\text{local}}, l_p)$-local constancy property holds for this model if

$$P\left(\left\|D \oslash \Omega\right\|_p < \epsilon\right) > 1 - \delta.$$

Here we need to assume that $\delta < 1$. The ϵ parameter controls the degree of local constancy and the δ parameter is indicative of our prior belief about the local constancy.

One can also think of defining local constancy by replacing Ω in sections 4.1 and 4.2 by the partial correlation matrix Π. Even though the definition

makes sense, the natural transition from the difference matrix used in [14] to our analogue D^a will not be consistent with it. The difference terms corresponding to partial correlation matrix will be

$$|\pi_{jk} - \pi_{jk'}| = \frac{1}{\sqrt{\omega_{jj}}} \left| \frac{\omega_{jk}}{\sqrt{\omega_{kk}}} - \frac{\omega_{jk'}}{\sqrt{\omega_{k'k'}}} \right|.$$

But, when we regress X_j on the rest of the nodes, then the local difference we aim to penalize is $\frac{1}{\omega_{jj}}|\omega_{jk} - \omega_{jk'}|$. In order to be consistent with our definition of neighborhood fused lasso, we need to replace $D \oslash \Omega$ by $(D\Omega_d) \oslash \Pi$, where $\Omega_d = (\text{diag}(\Omega))^{\frac{1}{2}}$. Assuming that the individual conditional variances are bounded away from zero (assumption A2 in section 6.1.1), this rescaling does not affect the local constancy and our penalty is consistent with the definition. Partial correlations have the advantage of all being on the same scale. This might in particular simplify the choice of ϵ.

5. Optimization method

First we briefly review the existing algorithms used for optimization of fused lasso type problems. Broadly speaking, there are two fundamental classes of algorithms used in this type of problems: (a) Solution path algorithms - which finds the entire solution for all values of the regularization parameters and (b) Approximation algorithms - which attempt to solve a large scale optimization given a fixed set of regularization parameters using first order approximations.

Friedman et al. [10] formulated the path-wise optimization method for the standard fused lasso signal approximation problem where the design matrix $X = I$. The algorithm was two-step and the final solution is obtained by soft-thresholding the total-variation norm penalized estimate obtained in the first step. The basic challenge in applying the coordinate descent algorithm to a fused lasso problem is the non-separability of the total variation penalty unlike usual lasso where the l_1-penalty is completely separable. So they used a modified coordinate descent approach where the *descent* step was followed by an additional fusion and smoothing step. However, this method works only for the total variation penalty and does not extend to fused lasso regression problems with generalized fusion penalty like our situation. Hoefling [13] proposed a path-wise optimization algorithm for generalized fusion penalties in fused lasso signal approximation problem. This algorithm uses the fact that when varying the penalty parameter, the corresponding sets of fused coefficients do only change at finitely many values of the penalty parameter. Tibshirani and Taylor [24] proposed another path algorithm for the fused lasso regression problem with a generalized fused lasso penalty by solving the dual optimization problem. However, the path algorithm they devised can be applied only when the design matrix is full rank and hence is not applicable for high dimensional problems. In order to resolve the problem when the matrix is not full rank, they proposed to add an infinitesimal perturbation $\epsilon\|\beta\|^2$. However, this does not solve the problem completely as a small ϵ leads to ill-conditioning and increasing number of

rows in the generalized fused penalty matrix causes inefficient solutions because of an increasing number of dual variables.

Approximation algorithms were developed to find efficient solutions to general fused lasso problem with a fixed set of penalty parameters, regardless of its rank and they usually adapt themselves to high dimensional problems quite easily. Usually these approximation algorithms are based on first order approximation type methods like gradient descent. Liu et al. [17] proposed the efficient fused lasso algorithm (EFLA) that solves standard fused lasso regression problem by replacing the quadratic error term in the optimization function by its first order Taylor expansion at an approximate solution followed by an additional quadratic regularization term. The approximate objective function has a fused lasso signal approximation form and can be solved by applying gradient descent on its dual which is a box-constrained quadratic program. Chen et al. [7] proposed the *smoothing proximal gradient* method to solve regression problem with structured penalties that closely resemble our objective function. The basic idea is to approximate the fused penalty term $\|D\beta\|_1$ by a smooth function and devise an iterative scheme for an approximate optimization. They use the smooth function

$$\tilde{\Omega}(\beta, t) = \max_{\|\alpha\|_\infty \leq 1} \left(\alpha^T D\beta - \frac{t}{2}\|\alpha\|_2^2 \right).$$

Convexity and continuous differentiability of this function follows from Nesterov [19]. One may now proceed using standard first order approximation approach like FISTA [2] which works like EFLA. However this process is computationally intensive and not a feasible approach in our case where we need to do a nodewise regression. Another algorithm for tackling a similar problem is known as *Split Bregman* algorithm [27]. It was proposed for standard fused lasso regression and later extended to generalized fused lasso regression. The SB algorithm is derived from augmented Lagrangian [12, 20], which adds quadratic penalty terms to the penalized objective function and alternatingly solves the primal and dual starting from an initial estimate. This is also computationally intensive unless one has a simple structure constraint on the parameters.

We propose a different algorithm to solve our optimization problem. We show that the neighborhood-fused lasso problem can be reparametrized into a standard lasso problem, thus simplifying the optimization procedure.

The neighborhood-fused LASSO estimate $\hat{\theta}^{a,\lambda,\mu}$ of θ^a can be written as

$$\begin{aligned}
\hat{\theta}^{a,\lambda,\mu} &= \text{argmin}_{\theta:\theta_a=0} \left(\frac{1}{n}||X_a - X\theta||^2 + \lambda||\theta||_1 + \mu||D^a\theta||_1 \right) \\
&= \text{argmin}_{\theta:\theta_a=0} \left(\frac{1}{n}||X_a - X\theta||^2 + \lambda \left(||\theta||_1 + \frac{\mu}{\lambda}||D^a\theta||_1 \right) \right) \\
&= \text{argmin}_{\theta:\theta_a=0} \left(\frac{1}{n}||X_a - X\theta||^2 + \lambda \left|\left| \begin{pmatrix} I \\ \frac{\mu}{\lambda}D^a \end{pmatrix} \theta \right|\right|_1 \right). \qquad (5.1)
\end{aligned}$$

Letting $G_a = \begin{pmatrix} I \\ \frac{\mu}{\lambda} D^a \end{pmatrix}$ and $\omega = G_a\theta$, we define

$$\hat{\omega}^{a,\lambda,\mu} = \text{argmin}_\omega \left(\frac{1}{n} ||X_a - XG_a^+\omega||^2 + \lambda||\omega||_1 \right), \tag{5.2}$$

where G_a^+ is the Moore-Penrose inverse of G_a. We find $\hat{\theta}^{a,\lambda,\mu}$ from the following relation (see lemma 5.1 for details)

$$\hat{\theta}^{a,\lambda,\mu} = G_a^+ \hat{\omega}^{a,\lambda,\mu}. \tag{5.3}$$

We note that since G_a has full column rank, G_a^+ has full row rank. Thus, the following lemma can be applied here.

Lemma 5.1. *Let*

$$\tilde{\beta} := \text{argmin}_{\beta \in \mathbb{R}^k} \|y - X\beta\|^2 + \lambda \|G\beta\|_1$$

and

$$\tilde{\omega} := \text{argmin}_\omega \|y - XG^+\omega\|^2 + \lambda \|\omega\|_1.$$

Also assume that G is of full column rank. Then

$$\tilde{\beta} = G^+\tilde{\omega}.$$

One interesting observation here is that although we started with the assumption that $\omega = G_a\theta \in \mathcal{C}(G_a)$, where $\mathcal{C}(G_a)$ denotes the column space of G_a, we did not optimize over $\mathcal{C}(G_a)$ but did the same over all ω. See the proof of lemma 5.1 for details. The heuristic idea behind this is that we find the minimizer on $\mathcal{C}(G_a)$ by projecting the global minimizer onto $\mathcal{C}(G_a)$ and the projection operator is given by $G_aG_a^+$. The objective function in (5.2) combines the two l_1 penalties into a single l_1 penalty and thus could be easily solved by any of the standard LASSO algorithms. The parameters λ and μ are chosen according to theorem 6.7. We show by several simulations that the proposed method performs better than Meinshausen - Bühlmann's method or graphical lasso in situations where local constancy holds.

6. Theoretical Properties of NFL

6.1. Model selection consistency and choice of penalty paramteters

From discussion in section 3, it can be seen that using neighborhood-fused lasso leads to efficient model selection when applied successively to all the nodes. In this section, we show that our proposed method leads to asymptotically consistent model selection similar to the procedure by Meinshausen and Bühlmann. The choice of the regularization parameters is crucial in such cases. Moreover, we show that our proposed choice of regularization parameters not only ensures convergence to the "true" model but the convergence is faster than the Meinhausen-Bühlmann's method when the underlying model is indeed locally constant. This property is also supported by our simulation results shown in section 7.

6.1.1. Assumptions

In order to prove consistency of our method for Gaussian graphical models, we need to work with the following assumptions. Assumptions [A1]-[A3] and [A5]-[A7] are as in Meinshausen and Bühlmann's (See [18], section 2.3). However we need the two additional assumptions [A4] and [A8] to deal with the local constancy.

[A1] **Dimensionality** : $\exists \gamma > 0$, such that $p(n) = O(n^{\gamma})$ as $n \to \infty$. Note that $\gamma > 1$ is included thus allowing $p \gg n$.

[A2] For all $a \in \Gamma(n)$ and $n \in \mathbb{N}$, $Var(X_a) = 1$. There exists $v^2 > 0$, so that for all $n \in \mathbb{N}$ and $a \in \Gamma(n)$, $Var(X_a | X_{\Gamma(n)\setminus\{a\}}) \geq v^2$. This means that all the conditional variances are bounded away from 0.

[A3] **Sparsity** : There exists some $0 \leq \kappa < 1$ so that $\max_{a\in\Gamma(n)} |\mathrm{ne}_a| = O(n^{\kappa})$ for $n \to \infty$.

[A4] **Local Neighborhood Sparsity**: The number of local neighbors also can grow at polynomial rate of n, i.e., $\exists \beta_0 \geq 0$ such that the maximum number of local neighbors of a node $\max_{a\in\Gamma(n)} |\mathrm{ne}_l(a)| = O(n^{\beta_0})$ for $n \to \infty$. For convenience we let $K > 0$ be such that

$$2 + \max_{a\in\Gamma(n)} |\mathrm{ne}_l(a)| < K n^{\beta_0}.$$

[A5] **l_1-Boundedness**: There exists some $\vartheta < \infty$ so that for all neighboring nodes $a, b \in \Gamma(n)$ and all $n \in \mathbb{N}$, $\|\theta^{a,ne_b\setminus\{a\}}\|_1 \leq \vartheta$.

[A6] **Magnitude of partial correlation** : There exists a constant $\delta > 0$ and some $\xi > \kappa$ with κ as in [A3], so that for every $(a, b) \in E$, $|\pi_{ab}| \geq \delta n^{-\frac{1-\xi}{2}+\beta_0}$ where π_{ab} denotes the partial correlation of X_a and X_b.

[A7] **Neighborhood Stability** : Define $S_a(b) = \sum_{k\in \mathrm{ne}_a} \mathrm{sgn}(\theta_k^{a,\mathrm{ne}_a})\theta_k^{b,\mathrm{ne}_a}$. There exists some $\delta_1 < 1$ so that for all $a, b \in \Gamma(n)$ with $b \notin \mathrm{ne}_a$, $|S_a(b)| < \delta_1$

[A8] **Local Neighborhood Stability**: Let $\mathcal{L}_a := \{k : ((D^a)' \mathrm{sgn}(D^a\theta^a))_k \neq 0\}$ and $T_a(b) := \sum_{k\in\mathcal{L}_a} \left[(D^a)' \mathrm{sgn}(D^a\theta^a)\right]_k \theta_k^{b,\mathcal{L}_a}$. There exists some $\delta_2 < 1$ so that for all $a, b \in \Gamma(n)$ such that $b \notin \mathcal{L}_a$, $|T_a(b)| < \delta_2 \|D^a_{\cdot b}\|_1$, where $D^a_{\cdot b}$ denotes denotes the b^{th} column of D^a.

Similar to Meinshausen and Bühlmann's interpretation, we can describe an intuitive condition which implies the last two assumptions. Define

$$\theta^a(\eta_1, \eta_2) = \mathrm{argmin}_{\theta:\theta_a=0} E(X_a - \sum_{k\in\Gamma(n)} \theta_k X_k)^2 + \eta_1\|\theta\|_1 + \eta_2\|D^a\theta^a\|_1.$$

According to the characterization of ne_a derived from (3.2), $\mathrm{ne}_a = \{k \in \Gamma(n) : \theta_k^a(0, 0) \neq 0\}$. One can think of a two-dimensional perturbation approach in which one tweaks the parameters η_1 and η_2 in a way that the perturbed neighborhood $\mathrm{ne}_a(\eta_1, \eta_2) = \{k \in \Gamma(n) : \theta_k^a(\eta_1, \eta_2) \neq 0\}$ is identical to original neighborhood $\mathrm{ne}_a(0, 0)$. The following proposition shows that the two assumptions of neighborhood stability are fulfilled under this situation. The terms $S_a(b)$ and $T_a(b)$ measure the sub-gradient of the lasso penalty and the neighborhood-fused

lasso penalty respectively. Having a small l_1 bound on them enforces the stability of the estimated coefficients, and hence stability of neighbors. See section 9 for proof of the proposition.

Proposition 6.1. *If there exists some* $\eta_1 > 0, \eta_2 > 0$ *such that* $\mathrm{ne}_a(\eta_1, 0) = \mathrm{ne}_a(0, \eta_2) = \mathrm{ne}_a(0, 0)$. *Then,* $|S_a(b)| \leq 1$ *and* $|T_a(b)| \leq \|D^a_{.b}\|_1$. *Moreover,* $\mathrm{ne}_a(\eta_1, \eta_2) = \mathrm{ne}_a(0, 0)$.

We start the presentation of the theoretical results with a lemma (proven in section 9) characterizing the minimizer of our objective function in terms of its subdifferential.

Lemma 6.2. *Given* $\theta \in \mathbb{R}^{p(n)}$, *let* $G(\theta)$ *be a* $p(n)$ *dimensional vector with elements* $G_b(\theta) = -\frac{2}{n}\langle X_a - X\theta, X_b\rangle$. *Define*

$$\mathcal{D}^a_b = [(D^a)'\mathrm{sgn}(D^a\theta)]_b \qquad \textit{and} \qquad \mathcal{L}^a(\theta) = \{b : \mathcal{D}^a_b \neq 0\}.$$

A vector $\hat{\theta}$ *is a solution to the fused LASSO problem described above iff*

$$\begin{aligned}
&G_b(\hat{\theta}) = \lambda\mathrm{sgn}(\hat{\theta}_b) + \mu\mathcal{D}^a_b && \textit{when } \hat{\theta}_b \neq 0, b \in \mathcal{L}^a(\hat{\theta}),\\
&\lambda\mathrm{sgn}(\hat{\theta}_b) - \mu\left\|D^a_{.b}\right\|_1 \leq G_b(\hat{\theta}) \leq \lambda\mathrm{sgn}(\theta_b) + \mu\left\|D^a_{.b}\right\|_1 && \textit{when } \theta_b \neq 0, b \notin \mathcal{L}^a(\hat{\theta}),\\
&-\lambda + \mu\mathcal{D}^a_b \leq G_b(\hat{\theta}) \leq \lambda + \mu\mathcal{D}^a_b && \textit{when } \theta_b = 0, b \in \mathcal{L}^a(\hat{\theta}),\\
&|G_b(\hat{\theta})| \leq \lambda + \mu\left\|D^a_{.b}\right\|_1 && \textit{otherwise.}
\end{aligned}$$

This lemma builds the foundation of several of the following results and will be used frequently to prove them. Sign consistency is one of the major properties that a model selection method should exhibit. Before we study the model selection consistency of our estimators, we show, in the following lemma that the neighborhood-fused lasso estimator is sign-consistent.

Lemma 6.3. *Let* $\hat{\theta}^{a,\lambda,\mu}$ *be defined for all* a. *Under the assumptions [A1]-[A7], it holds for some* c *that for all* a, $\mathbb{P}(\mathrm{sgn}(\hat{\theta}^{a,\lambda,\mu}_b) = \mathrm{sgn}(\theta^a_b)\forall b \in ne_a) = 1 - O(\exp(-cn^{\epsilon}))$.

Observe that this lemma, in turn, preserves the asymptotic equality of signs of local neighbors. If X_b and X'_b are local neighbors, then with high probability, when regressing X_a on the remaining variables, the coefficients of X_b and $X_{b'}$ will have the same sign. This is a direct consequence of local constancy of the estimated regression coefficients.

Our results show that, just like in Meinshausen and Bühlmann, a rate slower than $n^{-1/2}$ is necessary for the regularization parameters for consistent model selection in the high dimensional case where the dimension may increase as a polynomial in the sample size. Specifically if λ decays at $n^{-\frac{1-\varepsilon}{2}}$ and μ decays as $n^{-\frac{1-\varepsilon}{2}-\beta_0}$ where $0 < \beta_0 < \kappa < \varepsilon < \xi$ are as in assumptions [A1]-[A8], the estimated neighborhood is almost surely contained in the true neighborhood. Hence the type-I error probability goes to 0. This is formally stated in the following theorem (see section 9 for a proof).

Theorem 6.4. *Let assumptions [A1]-[A8] be fulfilled. Let the penalty parameters satisfy* $\lambda_n \sim d_1 n^{-\frac{1-\varepsilon}{2}}$ *and* $\mu_n \sim d_2 n^{-\frac{1-\varepsilon}{2}}$ *with some* $\beta_0 \geq 0$ *and* $0 < \kappa < \varepsilon < \xi$ *and* $d_1, d_2 > 0$. *There exists some* $c > 0$ *such that for all* $a \in \Gamma(n)$,

$$P(\hat{\text{ne}}_a^{\lambda,\mu} \subseteq \text{ne}_a) = 1 - O(\exp(-cn^{\varepsilon})).$$

The assumptions of neighborhood stability and local neighborhood stability are not redundant. The following proposition shows that one can not relax the assumptions A7 and A8.

Proposition 6.5. *If there exists some* $a, b \in \Gamma(n)$ *with* $b \notin \text{ne}_a$ *and* $|S_a(b)| > 1$ *and* $|T_a(b)| > \left\|D^a_{\cdot b}\right\|_1$, *then for* λ_n, μ_n *as in theorem 4.6,*

$$P\left(\hat{\text{ne}}_a^{\lambda,\mu} \subseteq \text{ne}_a\right) \to 0 \quad as \quad n \to \infty.$$

On the other hand, with the same sets of assumptions, one can show that the type II probability, i.e., probability of falsely identifying an edge as a potential connection exponentially goes to 0. This has been formally stated and proved in the following theorem (see section 9 for a proof).

Theorem 6.6. *With all the assumptions of theorem 4.6 and* λ, μ *as before,*

$$P\left(\text{ne}_a \subseteq \hat{\text{ne}}_a^{\lambda}\right) = 1 - O(\exp(-cn^{\varepsilon})).$$

On Model Selection in Graphs

It follows from the discussion so far that consistent estimation of nodes is possible using our approach. However, one of the original goals of our method is to estimate the entire underlying graphical model, which can be accomplished by combining the estimated neighborhoods in some way. Two different methods of combination have been proposed

$$\hat{E}^{\lambda,\mu,\vee} := \left\{(a,b) : a \in \hat{\text{ne}}_b^{\lambda,\mu} \vee b \in \hat{\text{ne}}_a^{\lambda,\mu}\right\} \quad \text{(union)};$$
$$\hat{E}^{\lambda,\mu,\wedge} := \left\{(a,b) : a \in \hat{\text{ne}}_b^{\lambda,\mu} \wedge b \in \hat{\text{ne}}_a^{\lambda,\mu}\right\} \quad \text{(intersection)}.$$

In our simulations, we combined them following the union method. It has been observed that the differences vanish asymptotically when a regular lasso is applied. Although we did not theoretically study this for our case but our experiments indicate that they do not vary much asymptotically.

Finite Sample Choices for Penalty Parameter

Theoretic exploration in asymptotic domain does not cast much light on the choices of regularization parameter in real life, finite sample problem. It is hard to ensure consistency or absolute containment like theorem 6.4 or 6.6. However, following the idea proposed by Meinshausen and Bühlmann, one can consider

the *connectivity component* of a node, which is defined as the set of nodes which are connected to it through a chain of edges. Generalizing the results from Meinshausen & Bühlmann, the following theorem shows (see proof in section 9) that the estimated connectivity component derived from neighborhood-fused lasso estimate will belong to the true connectivity component with probability $(1-\alpha)$, for any chosen level of $\alpha \in (0,1)$.

Theorem 6.7. *With all the assumptions [A1]-[A8], and the following choices of the penalty parameters,*

$$\lambda = \frac{\hat{\sigma}_a}{\sqrt{n}} \tilde{\Phi}^{-1}\left(\frac{\alpha}{2p(n)^2}\right)$$
$$\mu = \frac{\hat{\sigma}_a}{Kn^{\beta_0+1/2}} \tilde{\Phi}^{-1}\left(\frac{\alpha}{2p(n)^2}\right)$$

we have

$$P\left(\exists a \in \Gamma(n) : \hat{C}_a^{\lambda,\mu} \not\subseteq C_a\right) \leq \alpha$$

for all n. C_a and $\hat{C}_a^{\lambda,\mu}$ are the true and estimated connectivity components of a, K, β_0 are certain constants and $\tilde{\Phi} = 1 - \Phi$.

The choice of K and β_0 depends on the rate of growth of the local neighborhood with increasing sample size. In our simulations, we found that for models with constant dimension and increasing sample size, $K = 1$ and $\beta \in \left(0, \frac{1}{2}\right)$ works fine.

As shown in the simulations, we get a higher convergence speed as compared to Meinshausen-Bühlmann's method by applying a neighborhood-fused lasso penalty for nodewise regression when the underlying model exhibit local constancy. This can be proved using the following lemma (proof in section 9).

Lemma 6.8. *The upper bound of type I error probability in neighborhood fused lasso is smaller than that of the Meinshausen & Bühlmann procedure.*[1]

Lemma 6.8 does not specify how much the maximal false positive probability is reduced. This can be understood after going through its proof. Roughly speaking, the reduction in the upper bound is given by

$$\left[1 - e^{-\frac{1}{\sigma_*^2}\left(\frac{d_1 d_2}{2}(1-\delta_1)(1-\delta_2)n^{\beta_0} + \frac{d_2^2}{4}(1-\delta_2)^2 n^{2\beta_0}\right)n^{\epsilon}}\right] \cdot e^{-\frac{1}{\sigma_*^2}\left(\frac{d_1^2}{4}(1-\delta_1)^2 n^{\epsilon}\right)},$$

where $\sigma_*^2 = E(X_{a,i}^2 V_{a,i}^2)$, and δ_1, δ_2 are as in the proof of theorem 6.4. Essentially, it is shown in the proof that the neighborhood fused lasso reduces a dominant portion of the type 1 error (given by the second factor) by a fraction (given by the first factor) shown above. Looking at the proof of theorem 6.4, it is seen that the

[1] The lemma 6.8 delves deep into the precise constants of (9.1) in the proof of theorem 6.4 and shows that one could achieve smaller type 1 error probability with a finite sample size if one uses neighborhood fused lasso instead of usual lasso.

term $P\left(\left|2n^{-1}\langle X_a, V_b\rangle\right| \geq (1-\delta_1)\lambda + (1-\delta_2)B\mu\right)$ is the principal contributor to the probability of false positives. The corresponding term using usual lasso is $P\left(\left|2n^{-1}\langle X_a, V_b\rangle\right| \geq (1-\delta_1)\lambda\right)$. Lemma 6.8 makes use of this fact and it also shows that substantial reduction takes place when the local neighborhood grows. So, this method invariably performs better than the Meinshausen-Bühlmann's method with respect to a minimax criterion (in a sense that is minimizes the maximum probability of false positives), and the improvement is more when the local neighborhood grows faster. It also implies that it can perform as bad as usual lasso regression as the worst case scenario.

6.2. Compatibility and l_1 properties

In our attempt to Gaussian graphical model learning, we adopt the Meinshausen-Bühlmann approach, i.e., do a componentwise penalized regression. However, one should keep in mind that in the process of doing so, the design matrix (which is random here) changes at every iteration. In previous section we discussed the asymptotic model selection consistency of our estimator. In this section, we shall go beyond model selection and explore conditions under which our neighborhood-fused lasso estimate exhibit nice asymptotic l_1 properties. We shall carry out our theoretical analysis under the assumption that the linear regression model holds exactly, with some underlying "true" θ^0. Most of the theoretical results in this section have been derived in the light of discussions in [5].

We start with a quick recapitulation of the notation we are going to use here. If $X = (X_1, X_2, \cdots, X_p) \sim N(0, \Sigma)$, one assumes a node-wise regression model

$$X_a = X^a\theta_a + \epsilon_a \text{ for } a = 1, 2, \cdots, p, \tag{6.1}$$

where X_a denotes the a-th component, X^a denotes all the components except a and $\epsilon \sim N(0, \sigma^2 I_n)$ for some σ^2. Also assume that $\Sigma := ((\sigma_{ij}))_{i,j=1,\ldots,p}$ and $\Omega := \Sigma^{-1} = \left((\sigma^{ij})\right)_{i,j=1,\ldots,p}$ is the precision matrix. If E denotes the edge set in the conditional independence graph then $(i,j) \notin E \Leftrightarrow \sigma^{ij} = \sigma^{ji} = 0$. The design matrix X^a is obtained by deleting the a-th column of the data matrix.

In the high dimensional set up, where one generally has lesser number of samples that model dimension ($n < p$), we assume an inherent sparsity in the true θ^a. This sparsity is also ensured if we assume that the underlying conditional independence graph is sparse. Let us assume

$$S_a^0 = \{j : \theta_{a,j} \neq 0\}$$

and $s_a^0 = \text{card}(S_a^0)$. Since S_a^0 is not known, one needs a regularization penalty. Meinshausen and Bühlmann chose the l_1 penalty, i.e., the traditional lasso and got the estimate

$$\hat{\theta}_a^\lambda = \text{argmin}_{\theta_a}\left[\frac{1}{n}\left\|X_a - X^a\theta_a\right\|^2 + \lambda\left\|\theta_a\right\|_1\right].$$

In our situation, we have added another penalty term $\left\|D^a\theta_a\right\|_1$ along with the lasso penalty term. Hence, our estimator is given by

$$\hat{\theta}_a^{\lambda,\mu} = \text{argmin}_{\theta_a}\left[\frac{1}{n}\left\|X_a - X^a\theta_a\right\|^2 + \lambda\left\|\theta_a\right\|_1 + \mu\left\|D^a\theta_a\right\|_1\right].$$

Note that this new definition of our estimator is exactly same as what was defined in 3.3.

The Compatibility Condition for NFLasso

As mentioned in the earlier section, we shall develop our theory based on the assumption of a linear truth. We try to provide an upper bound on the prediction error. The following lemma forms the basis of our derivation.

Lemma 6.9. ***The basic inequality***

$$\frac{1}{n}\left\|X^a(\hat{\theta}_a^{\lambda,\mu} - \theta_a^0)\right\|^2 + \lambda\left\|\hat{\theta}_a^{\lambda,\mu}\right\|_1 + \mu\left\|D^a\hat{\theta}_a^{\lambda,\mu}\right\|_1 \leq \frac{2}{n}\epsilon_a' X^a(\hat{\theta}_a^{\lambda,\mu} - \theta_a^0) + \lambda\left\|\theta_a^0\right\|_1 + \mu\left\|D^a\theta_a^0\right\|_1.$$

It should be noted that the number of non zero elements in the column of the matrix D^a denotes the number of local neighbors that particular node has (except node a). Let us assume that the number of local neighbors is $O(n^{\beta_0})$. Fixing n, let us also assume that the number of local neighbors is bounded by B. Then, it follows from lemma 6.9 using $\left|\left\|x\right\|_1 - \left\|y\right\|_1\right| \leq \left\|x-y\right\|_1$ that

$$\frac{1}{n}\left\|X^a(\hat{\theta}_a^{\lambda,\mu} - \theta_a^0)\right\|^2 \leq \frac{2}{n}\epsilon_a' X^a(\hat{\theta}_a^{\lambda,\mu} - \theta_a^0) + (\lambda + B\mu)\left\|(\hat{\theta}_a^{\lambda,\mu} - \theta_a^0)\right\|_1.$$

The basic objective of using a penalty parameter is to overrule the empirical process term $\frac{2}{n}\epsilon_a' X^a(\hat{\theta}_a^{\lambda,\mu} - \theta_a^0)$. It can be easily seen that

$$\left|\frac{2}{n}\epsilon_a' X^a(\hat{\theta}_a^{\lambda,\mu} - \theta_a^0)\right| \leq \left(\frac{2}{n}\text{max}_{1\leq j\leq p}\left|\epsilon_a' X_j^a\right|\right)\left\|\hat{\theta}_a^{\lambda,\mu} - \theta_a^0\right\|_1.$$

Our goal is to choose a λ and a μ such that the probability that the empirical process term in the right hand side exceeds $\lambda + B\mu$ is small, so that with high probability the right hand side could be dominated by $(\lambda + B\mu)\left\|\hat{\theta}_a^{\lambda,\mu} - \theta_a^0\right\|_1$. To that effect, we define

$$\Lambda_a := \left\{\text{max}_{\substack{1\leq j\leq p \\ j\neq a}}\frac{2}{n}\left|\epsilon_a' X_j^a\right| \leq \lambda_0 + B\mu_0\right\}.$$

Our objective is to show that for some particular choice of λ_0 and μ_0, Λ_a has high probability. Now one should keep in mind that both ϵ_a and X_j^a are random here. One can make the valid assumption of their individual Gaussian law and independence. We formalize these notions in the following proposition.

Proposition 6.10. *Under the assumption of linear truth for the Gaussian graphical model, i.e., $X_a = X^a\theta_a^0 + \epsilon_a$ for $a = 1, 2, \cdots, p$, we have*

$$\epsilon_{a,i} \overset{\text{i.i.d.}}{\sim} N\left(0, \sigma_{aa} - \Sigma_{ab}\Sigma_{bb}^{-1}\Sigma_{ab}'\right) \quad \text{where} \quad \Sigma = \begin{pmatrix} \sigma_{aa} & \Sigma_{ab} \\ \Sigma_{ab}' & \Sigma_{bb} \end{pmatrix}.$$

The proof of this proposition is straightforward and hence skipped. The following results show that for suitable choice of λ_0 and μ_0, $P(\Lambda_a)$ is high.

Lemma 6.11. *Under the assumption that $\sigma_i = 1 \quad \forall i = 1, 2, \cdots, p$,*

$$P(\Lambda_a) \geq 1 - 2\exp\left[\log p - n\left(\frac{\lambda_0}{2} + 1 - \sqrt{\lambda_0 + B\mu_0 + 1}\right)\right],$$

In particular, with $\lambda_0 = \frac{2(t+\log p)}{n} > 0$ and $\mu_0 = \frac{2}{B}\sqrt{\frac{2}{n}(t + \log p)}$, then

$$P(\Lambda_a) \geq 1 - 2e^{-t}.$$

Corollary 6.12. *Assume that $\sigma_i = 1 \quad \forall i = 1, 2, \cdots, p$ and that $p = O(n^\gamma)$. Let the regularization parameters be*

$$\lambda = \frac{2\left(t^2 + \log p\right)}{n}, \quad \mu = \frac{1}{B}\sqrt{\frac{8(t^2 + \log p)}{n}}.$$

Then the following is true

$$P\left(\frac{1}{n}\left\|X^a(\hat{\theta}_a^{\lambda,\mu} - \theta_a^0)\right\|^2 \leq 2(\lambda + B\mu)\left\|\theta_a^0\right\|_1 + \mu B\left\|\hat{\theta}_a^{\lambda,\mu}\right\|_1\right) \geq 1 - e^{-t^2}.$$

The following lemma provides an upper bound on the estimation error.

Lemma 6.13. *Assume that $\sigma_i = 1 \quad \forall i = 1, 2, \cdots, p$ and that $p = O(n^\gamma)$. Let the regularization parameters be*

$$\lambda = \frac{2\left(t^2 + \log p\right)}{n}, \qquad \mu = \frac{1}{B}\sqrt{\frac{8(t^2 + \log p)}{n}}.$$

Also let $\delta_{\min}$ be the smallest non-zero singular value of X^a. Then we have

$$P\Big[\frac{1}{n}\|X^a(\hat{\theta}_a^{\lambda,\mu} - \theta_a^0)\|^2 \leq 2\left(\lambda + B\mu\left(1 + \frac{\sqrt{p}}{2}\right)\right)\|\theta_a^0\|_1 + \frac{npB\mu(\lambda + B\mu)}{\delta_{\min}}\Big] \geq 1 - e^{-t^2}.$$

Oracle Inequalities

We now try to derive some oracle inequalities which will provide l_1 bound for our neighborhood-fused lasso estimate. Following Bühlmann's notation, let us write, for an index set $S \subset \{1, 2, \cdots, p\}$,

$$\theta_{a,j,S} := \theta_{a,j}\mathbf{1}\{j \in S\},$$

$$\theta_{a,j,S^c} := \theta_{a,j}\mathbf{1}\{j \notin S\}.$$

We need some more notation. Split the matrix D^a as follows:

$$D^a = \begin{bmatrix} D^a_{S,S} & 0 \\ D^a_{S,0} & D^a_{0,S^c} \\ 0 & D^a_{S^c,S^c} \end{bmatrix},$$

where $D^a_{S,S}$ consists of all rows such that both the non zero terms belong to S. $D^a_{S,0}$ consists of all rows and columns such that exactly one of the non-zero term in that row belongs to S. D^a_{0,S^c} consists of all rows and columns such that exactly one of the non-zero term in that row belongs to S^c. $D^a_{S^c,S^c}$ consists of all rows such that both the non-zero terms belong to S.

Lemma 6.14. *On Λ_a, if we choose $\lambda \geq 2\lambda_0$ and $\mu \geq 2\mu_0$, we have*

$$\frac{2}{n}\left\|X^a\left(\hat{\theta}^{\lambda,\mu}_a - \theta^0_a\right)\right\|^2 + (\lambda - 3B\mu)\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0^c}\right\|_1 \quad \leq \quad (3\lambda + 5B\mu)\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0} - \theta^0_{a,S_0}\right\|_1.$$

It can be easily verified that if $\lambda \geq \left(3 + \frac{14}{\Delta}\right) B\mu$ for some $\Delta > 0$, we have

$$\lambda - 3B\mu \geq \frac{1}{3+\Delta}(3\lambda + 5B\mu).$$

This helps us to simplify the consequences lemma 6.14, which implies that

$$(\lambda - 3B\mu)\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0^c}\right\|_1 \quad \leq \quad (3\lambda + 5B\mu)\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0} - \theta^0_{a,S_0}\right\|_1.$$

Combining with the aforementioned condition, we get

$$\frac{3\lambda + 5B\mu}{3+\Delta}\|\hat{\theta}^{\lambda,\mu}_{a,S_0^c}\|_1 \;\leq\; (\lambda - 3B\mu)\|\hat{\theta}^{\lambda,\mu}_{a,S_0^c}\|_1 \;\leq\; (3\lambda + 5B\mu)\|\hat{\theta}^{\lambda,\mu}_{a,S_0} - \theta^0_{a,S_0}\|_1,$$

or,

$$\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0^c}\right\|_1 \quad \leq \quad (3+\Delta)\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0} - \theta^0_{a,S_0}\right\|_1.$$

A standard way to relate the l_1 penalty $\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0} - \theta^0_{a,S_0}\right\|_1$ to an l_2 penalty is to use the Cauchy-Schwarz inequality

$$\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0} - \theta^0_{a,S_0}\right\|_1 \leq \sqrt{s_0}\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0} - \theta^0_{a,S_0}\right\|,$$

and subsequently relate it to the l_2 penalty on the left hand side

$$\frac{2}{n}\left\|X^a\left(\hat{\theta}^{\lambda,\mu}_a - \theta^0_a\right)\right\|^2 = \frac{2}{n}\left(\hat{\theta}^{\lambda,\mu}_a - \theta^0_a\right)' X^{a'} X^a \left(\hat{\theta}^{\lambda,\mu}_a - \theta^0_a\right)$$

in the following manner

$$\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0} - \theta^0_{a,S_0}\right\|^2 \leq \frac{\left(\hat{\theta}^{\lambda,\mu}_a - \theta^0_a\right)' X^{a'} X^a \left(\hat{\theta}^{\lambda,\mu}_a - \theta^0_a\right)}{\phi^2_{0,a}},$$

where $\phi_{0,a} > 0$ is some constant. However, $\hat{\theta}_a^{\lambda,\mu}$ being random, this condition can not hold unanimously for all $\theta \in \mathbb{R}^p$. Bühlmann provided a *compatibility condition* so that this is true. In our situation, we found a similar condition like them to carry out further analysis. It should be noted that our compatibility condition is weaker than the compatibility condition Bühlmann provided. However, we need to work under the assumption that $\lambda \geq \left(3 + \frac{14}{\Delta}\right)$. The definition of this *compatibility condition* is provided below.

Definition 6.15.(Compatibility Condition) We say that a collection of nodes S_0 satsifies the Δ-compatibility condition if for all θ satisfying

$$\left\|\theta_{S_0^c}\right\|_1 \leq (3+\Delta)\left\|\theta_{S_0}\right\|_1$$

we have

$$\left\|\theta_{S_0}\right\|_1^2 \leq \frac{s_0\left(\theta' X^{a'} X^a \theta\right)}{n\phi_{0,a}^2}.$$

Following Bickel et al. [3], we shall refer to $\phi_{0,a}^2$ as the *restricted eigenvalue*. Here we provide a detailed explanation of this phenomenon. Observe that the compatibility condition could be re-written as

$$\phi_{0,a}^2 \leq \inf_{\theta:\left\|\theta_{S_0^c}\right\|_1 \leq (3+\Delta)\left\|\theta_{S_0}\right\|_1} \frac{s_0\left(\frac{1}{n}\theta' X^{a'} X^a \theta\right)}{\left\|\theta_{S_0}\right\|_1^2},$$

so that $\phi_{0,a}$ could be taken as the square root of the infimum assumed by the right hand side. Observing further that

$$\begin{aligned}
&\inf_{\theta:\left\|\theta_{S_0^c}\right\|_1 \leq (3+\Delta)\left\|\theta_{S_0}\right\|_1} \frac{s_0\left(\frac{1}{n}\theta' X^{a'} X^a \theta\right)}{\left\|\theta_{S_0}\right\|_1^2} \geq \inf_\theta \frac{s_0\left(\frac{1}{n}\theta' X^{a'} X^a \theta\right)}{\left\|\theta_{S_0}\right\|_1^2} \\
&\geq \inf_\theta \frac{s_0\left(\frac{1}{n}\theta' X^{a'} X^a \theta\right)}{s_0\left\|\theta_{S_0}\right\|^2} \geq \inf_\theta \frac{\left(\frac{1}{n}\theta' X^{a'} X^a \theta\right)}{\left\|\theta\right\|^2} = \lambda_{\min}\left(\frac{1}{n} X^{a'} X^a\right),
\end{aligned}$$

$\phi_{0,a}^2$ can be assumed to be the minimum eigenvalue over the restricted set $\{\theta : \left\|\theta_{S_0^c}\right\|_1 \leq (3+\Delta)\left\|\theta_{S_0}\right\|_1\}$. With this compatibility condition imposed, we derive the following oracle inequality

Theorem 6.16. *Assume that the compatibility condition (from definition 6.15) holds for some $\Delta > 0$. Then on Λ_a, for $\lambda \geq 2\lambda_0$, $\mu \geq 2\mu_0$ and $\lambda \geq \left(3 + \frac{14}{\Delta}\right) B\mu$, we have*

$$\frac{1}{n}\left\|X^a\left(\hat{\theta}_a^{\lambda,\mu} - \theta_a^0\right)\right\|^2 + (\lambda - 3B\mu)\left\|\hat{\theta}_a^{\lambda,\mu} - \theta_a^0\right\|_1 \quad \leq \quad \frac{s_0(2\lambda + B\mu)^2}{\phi_{0,a}^2}.$$

Combining theorem 6.16 and lemma 6.11, we get,

Theorem 6.17. *Let*

$$\lambda = \frac{4(t+\log p)}{n}, \quad \mu = \frac{4}{B}\sqrt{\frac{2}{n}(t+\log p)} \quad \textit{and} \quad t \geq 2n\left(3+\frac{14}{\Delta}\right)^2,$$

and assume that $\sigma_i = 1,\ i = 1, 2, \cdots, p$. *Then with probability* $1 - e^{-t}$,

$$\frac{1}{n}\left\|X^a\left(\hat{\theta}_a^{\lambda,\mu} - \theta_a^0\right)\right\|^2 + (\lambda - 3B\mu)\left\|\hat{\theta}_a^{\lambda,\mu} - \theta_a^0\right\|_1 \quad \leq \quad \frac{s_0(2\lambda + B\mu)^2}{\phi_{0,a}^2}.$$

Theorem 6.17 provides a unified way to deal with both the squared estimation error and the absolute error of the estimated parameters in neighborhood fused lasso regression. It can be seen that for increasing n, with probability increasing to 1, the absolute difference of our NFL estimator from underlying truth is bounded by

$$\frac{s_0(2\lambda + B\mu)^2}{\phi_{0,a}^2(\lambda - 3B\mu)}.$$

For large values of n, this is approximately equal to

$$\frac{s_0\left(\frac{8t}{n} + 4\sqrt{\frac{2t}{n}}\right)^2}{\phi_{0,a}^2\left(\frac{4t}{n} - 12\sqrt{\frac{2t}{n}}\right)},$$

which is a positive and decreasing function of $\frac{t}{n}$ if $\frac{t}{n} > 18$. This is automatically satisfied for all $\Delta > 0$. So, it can be easily seen, using the assumption $t \geq 2n\left(3 + \frac{14}{\Delta}\right)^2$, that the above quantity is bounded by

$$\frac{224 s_0(\Delta + 4)^2(3\Delta + 14)}{\phi_{0,a}^2 \Delta^2}.$$

7. Simulations

7.1. Simulation 1

In this simulation we repeat the exact scenario presented in Honorio's first simulation setting. The Gaussian graphical model consists of 9 variables as shown in figure 7.1. It deals with both local and non-local interactions. Our method is compared to Meinshausen-Bühlmann's method and graphical lasso. We run our simulation for 4 different sample sizes $n = 4, 50, 100$ & 400. For each sample size, we run the simulation 50 times and estimate the neighborhood for each iteration. We construct a weighted graph with edge weight corresponding to the frequency of its occurrence in all of those 50 iterations.

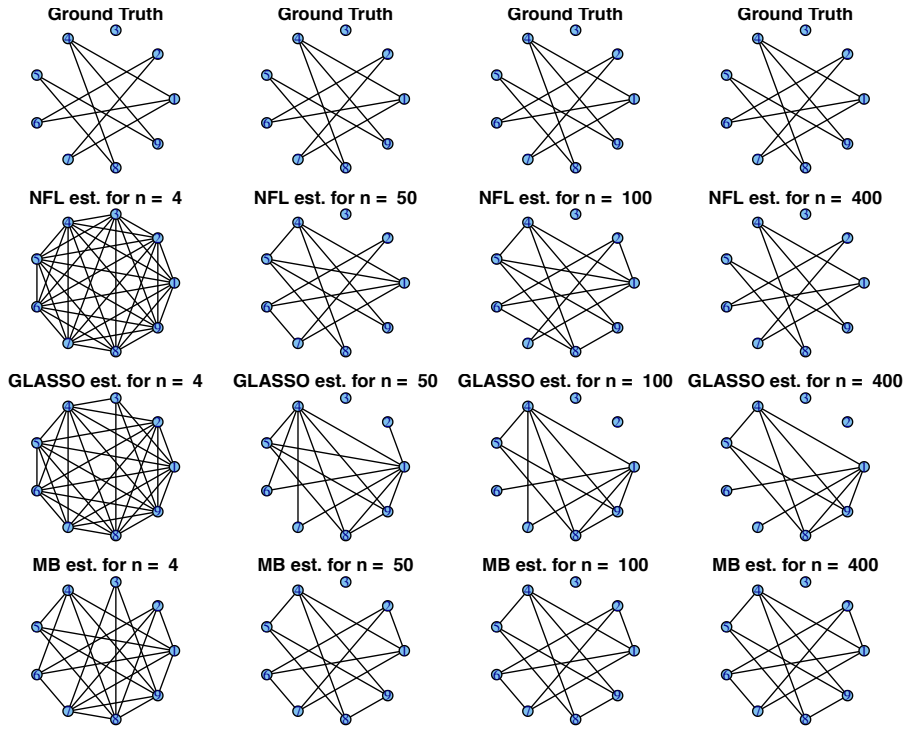
Ground Truth
Ground Truth
Ground Truth
Ground Truth
NFL est. for n = 4
NFL est. for n = 50
NFL est. for n = 100
NFL est. for n = 400
GLASSO est. for n = 4
GLASSO est. for n = 50
GLASSO est. for n = 100
GLASSO est. for n = 400
MB est. for n = 4
MB est. for n = 50
MB est. for n = 100
MB est. for n = 400

FIG 7.1. *Comparison of NFL with GLASSO and Meinshausen-Bühlmann estimates in section 7.1*

7.2. Simulation 2

In this simulation study we take a 50 dimensional normal random vector with zero mean. The diagonals of the precision matrix are all 1 and all the nonzero off-diagonal entries are 0.2. The conditional (in)dependence graph of

this vector consists of both spatial (local) and non-spatial neighbors where the local neighborhood structure is linear (one dimensional lattice). There are two groups of distant neighbors. We generate n i.i.d. samples from the corresponding normal distribution. We run the simulation for $n = 10, 25, 50, 100, 500, 1000$. We try to reconstruct the conditional (in)dependence graph from the data using Graphical LASSO (we use an oracle version of graphical lasso where the choice of penalty parameter is contingent on the actual number of edges in the true model), Meinshausen - Bühlmann's coordinate-wise LASSO and our coordinate-wise generalized Fused LASSO approach. For each n, we run the simulation 50 times and calculate the number of correctly identified edges and that of falsely identified edges for each iteration. The mean and standard deviations are shown in the table 7.1 and table 7.2. Figure 7.2 shows the relative performance of the competing methods for different sample sizes. Sample size increases from top to bottom. In the figure we include comparison for two additional sample values. They are $n = 5000$ and $n = 10000$ respectively.

Method	Parameters	$n = 10$		$n = 25$		$n = 50$	
		fp	tp	fp	tp	fp	tp
GL	$\rho = 0.00$	NA	NA	NA	NA	NA	NA
	$\rho = 0.05$	239.78(8.94)	24.98(4.13)	352.79(9.89)	40.33(3.86)	364.65(14.77)	52.08(3.79)
	$\rho = 0.10$	217.71(8.73)	23.20(3.83)	273.30(11.22)	36.10(3.57)	248.21(11.54)	45.58(3.52)
	$\rho = 0.15$	195.75(9.85)	21.78(3.68)	210.74(11.01)	31.92(3.88)	165.44(9.21)	38.640(3.70)
	$\rho = 0.20$	175.32(10.02)	20.60(3.25)	161.56(10.49)	27.88(3.45)	104.28(8.24)	31.64(3.75)
	$\rho = 0.25$	156.95(11.18)	18.86(3.28)	121.98(10.66)	23.42(3.84)	59.92(6.64)	24.740(4.13)
	$\rho = 0.30$	138.76(10.79)	17.32(3.42)	89.40(10.55)	19.34(3.51)	33.08(5.67)	17.53(3.15)
MB		0(0)	0(0)	0(0)	0(0)	0(0)	0(0)
NFL		51.62(5.74)	7.54(2.69)	37.32(4.84)	10.12(2.69)	47.20(6.05)	20.18(4.02)

TABLE 7.1
Comparison of False Positives and True Positives

Method	Parameters	$n = 100$		$n = 500$		$n = 1000$	
		fp	tp	fp	tp	fp	tp
GL	$\rho = 0.00$	579.72(22.16)	64.06(3.03)	580.64(20.99)	69(0)	584.7(24.48)	69(0)
	$\rho = 0.05$	336.84(15.68)	60.34(2.79)	166.10(8.39)	68.24(0.85)	83.16(7.06)	68.9(0.36)
	$\rho = 0.10$	187.14(11.43)	52.46(2.97)	23.90(4.49)	61.26(1.66)	2.7(1.76)	62.74(1.64)
	$\rho = 0.15$	94.48(9.24)	42.50(3.59)	1.5(1.04)	49.62(1.71)	0.02(0.14)	50.28(1.84)
	$\rho = 0.20$	38.80(6.33)	32.48(3.25)	0.04(0.20)	33.62(2.92)	0(0)	35.14(3.09)
	$\rho = 0.25$	14.30(4.45)	21.76(3.32)	0(0)	15.12(3.99)	0(0)	12.86(2.67)
	$\rho = 0.30$	3.96(2.08)	13.32(2.87)	0(0)	3.6(1.96)	0(0)	1.12(0.87)
MB		0(0)	0(0)	0(0)	0(0)	0(0)	0.95(0.92)
NFL		99.92(7.08)	41.10(3.73)	2.44(1.40)	49.86(2.42)	0.02(0.14)	49.7(2.25)

TABLE 7.2
Comparison of False Positives and True Positives

It is quite evident from the two simulation that our method converges to

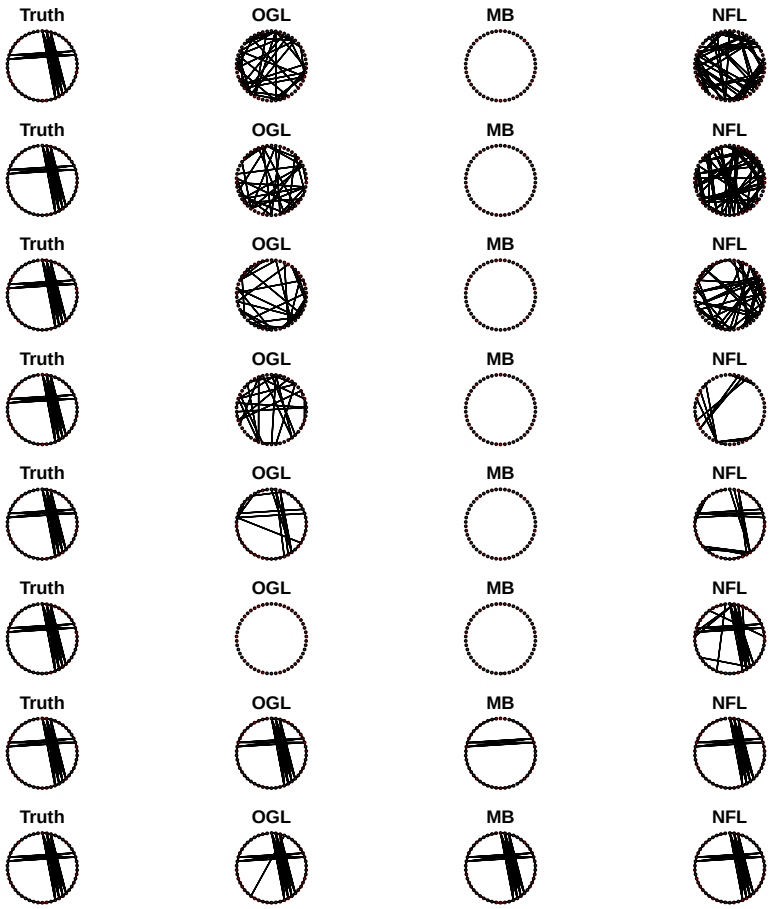

FIG 7.2. *Comparison of NFL with GLASSO and Meinshausen-Bühlmann estimate in section 7.2, sample sizes from top to bottom are 10, 25, 50, 100, 500, 1000, 5000, 10000*

the true model faster than Meinshausen-Bühlmann's method. It is also shown theoretically in lemma 6.8.

8. Discussion and Future Works

In this paper, we have successfully extended the Meinshausen-Bühlmann approach of model selection in Gaussian graphical models where the assumption of local constancy holds. We also provided a rigorous and generalized definition of local constancy and used it to propose neighborhood-fused lasso (NFL), an algorithm to solve the problem by penalizing the differences of local neighbors given by a pre-determined graph $\mathcal{G}_{\text{local}}$. We have used a generalized version of fused lasso in order to accomplish our objective. Our algorithm reduces the fused lasso problem into a regular lasso problem for a given set of regularizing parameters and thereby fast solutions are readily available. We were able to provide data dependent choices for the tuning parameters in order to establish asymptotic consistency in model selection. We substantiated our theoretical discussion on asymptotic model selection consistency with simulations that furnished desired results. We also proved, both theoretically and experimentally, that incorporating local constancy ensures faster convergence to the underlying truth, meaning that similar accuracy is achieved with smaller sample size. This phenomenon is somewhat similar to what Buhl [4] and Uhler [25] observed while investigating the existence of MLE's in Gaussian models with certain geometric structures. They were able to prove the existence of MLE of a higher dimensional model with relatively smaller sample size. Our findings are somewhat reminiscent of these results. We also discussed about the compatibility issue while doing regression with local constancy and were able to derive sufficient conditions under which l_1 boundedness of estimated parameters is ensured. On top of that we also provided an upper bound for the quadratic prediction error.

There are many extensions possible for future research. Firstly, instead of using the pre-defined local graph $\mathcal{G}_{\text{local}}$, one can think of alternative ways to grow or shrink the local neighborhood adaptively. Secondly, as discussed in section 4, one can think of devising an algorithm that penalizes the differences of partial correlations instead of the elements of the precision matrix. Thirdly, (because of the compatibility issue) regular lasso or fused lasso might fail to perform up to the mark if the model is ill-conditioned or the compatibility assumptions do not hold. One can think of using adaptive lasso or modify it accordingly in order to incorporate the local constancy property. Some preliminary simulations in this direction look promising. Thirdly, a Bayesian approach could be pursued where one can start with some prior distribution on the inverse covariance matrix respecting the local constancy property and use the Bayesian definition of local constancy in this paper (see 4.2) to come up with a novel alternative method.

Appendix.

9. Proofs of Theoretical Results

Proof of lemma 5.1. Let $G^+\tilde{\omega} = \beta_0$. We will show that $\beta_0 = \tilde{\beta}$. We know from definition of $\tilde{\omega}$ that $\forall \omega$

$$\left\|y - XG^+\tilde{\omega}\right\|^2 + \lambda\left\|\tilde{\omega}\right\|_1 \leq \left\|y - XG^+\omega\right\|^2 + \lambda\left\|\omega\right\|_1.$$

Since the objective function is convex, the minimizer in $\mathcal{C}(G)$ is obtained by projecting the grand minimizer onto $\mathcal{C}(G)$. Hence,

$$\hat{\omega} := \text{argmin}_{\omega \in \mathcal{C}(G)} \left\|y - XG^+\omega\right\|^2 + \lambda\left\|\omega\right\|_1 = GG^+\tilde{\omega}.$$

Therefore, for any $\omega \in \mathcal{C}(G)$, we have

$$\begin{aligned}
&\left\|y - XG^+\hat{\omega}\right\|^2 + \lambda\left\|\hat{\omega}\right\|_1 \leq \left\|y - XG^+\omega\right\|^2 + \lambda\left\|\omega\right\|_1 \\
\Rightarrow &\left\|y - XG^+GG^+\tilde{\omega}\right\|^2 + \lambda\left\|GG^+\tilde{\omega}\right\|_1 \leq \left\|y - XG^+\omega\right\|^2 + \lambda\left\|\omega\right\|_1 \\
\Rightarrow &\left\|y - X\beta_0\right\|^2 + \lambda\left\|G\beta_0\right\|_1 \leq \left\|y - XG^+G\theta\right\|^2 + \lambda\left\|G\theta\right\|_1 \quad \forall\theta.
\end{aligned}$$

Since G is full column rank, $G^+G = I$. Hence,

$$\left\|y - X\beta_0\right\|^2 + \lambda\left\|G\beta_0\right\|_1 \leq \left\|y - X\theta\right\|^2 + \lambda\left\|G\theta\right\|_1 \quad \forall\theta.$$

Therefore we get that

$$\beta_0 = \text{argmin}_{\beta \in \mathcal{C}(G^+) = \mathbb{R}^k} \left\|y - X\beta\right\|^2 + \lambda\left\|G\beta\right\|_1.$$

□

Proof of proposition 6.1. The proof can be found in Supplement A. □

Proof of lemma 6.2. The subdifferential of $\frac{1}{n}\left\|X_a - X\theta\right\|_2^2 + \lambda\left\|\theta\right\|_1 + \mu\left\|D^a\theta\right\|_1$ is given by $\{G(\theta) + \lambda e_1 + \mu e_2 : e_1 \in S_1, e_2 \in S_2\}$ where $S_1 = \{e \in \mathbb{R}^{p(n)} : e_b = \text{sgn}(\theta_b) \quad if \quad \theta_b \neq 0 \quad and \quad e_b \in [-1,1] \quad if \quad \theta_b = 0\}$ and $S_2 = \{e \in \mathbb{R}^{p(n)} : e_b = \alpha_b \quad if \quad \alpha_b \neq 0 \quad and \quad e_b \in [-1,1] \quad if \quad \alpha_b = 0$ where $\alpha = D'^a\text{sgn}(D^a\theta)\}$. Observe that $|(D'^a\text{sgn}(D^a\theta))_b| \leq \left\|D^a_{\cdot b}\right\|_1$ where $D^a_{\cdot b}$ denotes the b^{th} column of the matrix D^a. This is same as the total number of local neighbors of b except a. The lemma follows. □

Proof of lemma 6.3. This proof can be found in Supplement A. □

Proof of theorem 6.4. The event $\hat{\text{ne}}_a^{\lambda,\mu} \nsubseteq \text{ne}_a$ is equivalent to the event that there exists some node $b \in \Gamma(n) \setminus \text{cl}_a$ in the set of non-neighbors of node a such that the estimated coefficient $\hat{\theta}_b^{a,\lambda,\mu}$ is not zero. Thus,

$$P\left(\hat{\text{ne}}_a^{\lambda,\mu} \subseteq \text{ne}_a\right) = 1 - P\left(\exists b \in \Gamma(n) \setminus \text{cl}_a : \hat{\theta}_b^{a,\lambda,\mu} \neq 0\right).$$

Let $\mathcal{E}$ be the event that

$$\max_{k \in \Gamma(n)\setminus \text{cl}_a} \left| G_k\left(\hat{\theta}^{a,\text{ne}_a,\mathcal{B},\lambda,\mu}\right)\right| < \lambda + B\mu.$$

Conditional on $\mathcal{E}$, it follows from Meinshausen-Bühlmann's discussion that $\hat{\theta}^{a,\text{ne}_a,\mathcal{B},\lambda,\mu}$ is also a solution to the fused LASSO problem with $\mathcal{A} = \Gamma(n) \setminus \{a\}$. As $\hat{\theta}_b^{a,\text{ne}_a,\mathcal{B},\lambda,\mu} = 0$ for all $b \in \Gamma(n)\setminus \text{cl}_a$, it follows that $\hat{\theta}_b^{a,\lambda,\mu} = 0$ for all $b \in \Gamma(n)\setminus \text{cl}_a$. By 6.2 ,

$$P\left(\exists b \in \Gamma(n) \setminus \text{cl}_a : \hat{\theta}_b^{a,\lambda,\mu} \neq 0\right) \leq P\left(\max_{k \in \Gamma(n)\setminus \text{cl}_a} \left| G_k\left(\hat{\theta}^{a,\text{ne}_a,\mathcal{B},\lambda,\mu}\right)\right| \geq \lambda + B\mu\right).$$

It suffices to show that there exists a constant $c > 0$ so that for all $b \in \Gamma(n) \setminus \text{cl}_a$,

$$P\left(\left|G_b\left(\hat{\theta}^{a,\text{ne}_a,\lambda,\mu}\right)\right| \geq \lambda + B\mu\right) = O(\exp(-cn^{\varepsilon})).$$

Writing $X_b = \sum_{m\in \text{ne}_a} \theta_m^{b,\text{ne}_a} X_m + V_b$ for any $b \in \Gamma(n) \setminus \text{cl}_a$ where $V_b \sim N(0, \sigma_b^2)$ for some $\sigma_b^2 \leq 1$ and V_b is independent of $\{X_m; m \in \text{cl}_a\}$. Hence,

$$G_b\left(\hat{\theta}^{a,\text{ne}_a,\lambda,\mu}\right) = -2n^{-1} \sum_{m\in \text{ne}_a} \Big[\theta_m^{b,\text{ne}_a} \langle X_a - X\hat{\theta}^{a,\text{ne}_a,\lambda,\mu}, X_m\rangle - 2n^{-1}\langle X_a - X\hat{\theta}^{a,\text{ne}_a,\lambda,\mu}, V_b\rangle\Big]$$

By lemma 6.3, there exists a $c > 0$ so that with probability $1 - O(\exp(-cn^{\varepsilon}))$,

$$\text{sgn}\left(\hat{\theta}_k^{a,\text{ne}_a,\lambda,\mu}\right) = \text{sgn}\left(\theta_k^{a,\text{ne}_a}\right) \ \forall k \in \text{ne}_a.$$

In this case, it holds by lemma 6.2 that

$$\begin{aligned}\left|2n^{-1} \sum_{m\in \text{ne}_a} \theta_m^{b,\text{ne}_a} \langle X_a - X\hat{\theta}^{a,\text{ne}_a,\lambda,\mu}, X_m\rangle\right| &\leq \left|\sum_{m\in \text{ne}_a} \text{sgn}\left(\theta_m^{a,\text{ne}_a}\right) \theta_m^{b,\text{ne}_a}\lambda\right| \\ &\quad + \left|\sum_{m\in \text{ne}_a} [D'^a \text{sgn}(D^a\theta)]_m \theta_m^{b,\text{ne}_a}\mu\right| \\ &\leq \delta_1\lambda + \delta_2 B\mu\end{aligned}$$

The absolute value of the coefficient G_b is hence bounded by

$$\left|G_b\left(\hat{\theta}^{a,\text{ne}_a,\lambda,\mu}\right)\right| \leq \delta_1\lambda + \delta_2 B\mu + \left|2n^{-1}\langle X_a - X\hat{\theta}^{a,\text{ne}_a,\lambda,\mu}, V_b\rangle\right|.$$

with probability $1 - O(\exp(-cn^{\varepsilon}))$. Conditional on X_{cl_a}, the random variable $\langle X_a - X\hat{\theta}_k^{a,\text{ne}_a,\lambda,\mu}, V_b\rangle$ is normally distributed with mean 0 and variance $\sigma_b^2\|X_a - X\hat{\theta}_k^{a,\text{ne}_a,\lambda,\mu}\|^2$. By definition of $\hat{\theta}^{a,\text{ne}_a,\lambda,\mu}$,

$$\|X_a - X\hat{\theta}_k^{a,\text{ne}_a,\lambda,\mu}\|^2 \leq \|X_a\|^2$$

Since $\sigma_b^2 \leq 1$, $2n^{-1}\langle X_a - X\hat{\theta}_k^{a,\mathrm{ne}_a,\lambda,\mu}, V_b\rangle$ is stochastically smaller than or equal to $\left|2n^{-1}\langle X_a, V_b\rangle\right|$. It remains to show that for some $c > 0$ and some $0 < \delta_1, \delta_2 < 1$,

$$P\left(\left|2n^{-1}\langle X_a, V_b\rangle\right| \geq (1-\delta_1)\lambda + (1-\delta_2)B\mu\right) = O(\exp(-cn^{\varepsilon})).$$

If X_a and V_b are independent, $E(X_aV_b) = 0$. Using Gaussianity and bounded variance of both X_a and V_b, we obtain existence of some $g < \infty$ such that $E\left(\exp(|X_aV_b|)\right) < g$. Hence using Bernstein inequality and boundedness of λ and μ, it holds for some $c > 0$ that for all $b \in \mathrm{ne}_a$ (see lemma 6.8 for detailed proof),

$$P\left(\left|2n^{-1}\langle X_a, V_b\rangle\right| \geq (1-\delta_1)\lambda + (1-\delta_2)B\mu\right) = O(\exp(-cn^{\varepsilon})) \tag{9.1}$$

which completes the proof. □

Proof of proposition 6.5. We follow the proof of theorem 6.4 and claim that with the above assumptions, for all a, b with $b \notin \mathrm{ne}_a$,

$$P\left(|G_b\left(\hat{\theta}^{a,\mathrm{ne}_a,\lambda,\mu}\right)| > \lambda + B\mu\right) \to 1 \text{ for } n \to \infty$$

Following similar arguments afterwards, it can be concluded that with some $\delta_1 > 1$ and $\delta_2 > 1$ as $n \to \infty$

$$P\left(|G_b\left(\hat{\theta}^{a,\mathrm{ne}_a,\lambda,\mu}\right)| \geq \delta_1\lambda + \delta_2 B\mu - |2n^{-1}\langle X_a - X\hat{\theta}^{a,\mathrm{ne}_a,\lambda,\mu}, V_b\rangle|\right) \to 1.$$

It holds for the third term that for any $g_1 > 0, g_2 > 0$,

$$P\left(|2n^{-1}\langle X_a - X\hat{\theta}^{a,\mathrm{ne}_a,\lambda,\mu}, V_b\rangle| > g_1\lambda + g_2 B\mu\right) \to 0 \text{ as } n \to \infty.$$

which combined with the previous result proves the proposition. □

Proof of theorem 6.6. Observe that

$$P\left(\mathrm{ne}_a \subseteq \hat{\mathrm{ne}}_a^{\lambda}\right) = 1 - P\left(\exists b \in \mathrm{ne}_a : \hat{\theta}_b^{a,\lambda,\mu} = 0\right).$$

Let $\mathcal{E}$ be the event

$$\max_{k \in \Gamma(n)\setminus \mathrm{cl}_a} |G_k\left(\hat{\theta}^{a,\mathrm{ne}_a,\lambda,\mu}\right)| < \lambda + B\mu.$$

On $\mathcal{E}$, following similar arguments as before, we can conclude that $\hat{\theta}^{a,\mathrm{ne}_a,\lambda,\mu} = \hat{\theta}^{a,\lambda,\mu}$. Therefore

$$P\left(\exists b \in \mathrm{ne}_a : \hat{\theta}_b^{a,\lambda,\mu} = 0\right) \leq P\left(b \in \mathrm{ne}_a : \hat{\theta}^{a,\mathrm{ne}_a,\lambda,\mu} = 0\right) + P(\mathcal{E}^c).$$

It follows from the proof of Theorem 4.6 that there exists some $c > 0$ so that $P(\mathcal{E}^c) = O(\exp(-cn^{\varepsilon}))$. Using Bonferronis inequality, it hence remains to show that there exists some $c > 0$ so that for all $b \in \mathrm{ne}_a$,

$$P\left(\hat{\theta}^{a,\mathrm{ne}_a,\lambda,\mu} = 0\right) = O(\exp(-cn^{\varepsilon})),$$

which follows from lemma 6.3. □

Proof of theorem 6.7. $\hat{C}_a^{\lambda,\mu} \nsubseteq C_a$ implies the existence of an edge in the estimated neighborhood that connects two nodes in two different connectivity components of the true underlying graph. Hence,

$$P\left(\exists b \in \Gamma(n) : b \in \hat{\text{ne}}_a^{\lambda,\mu}\right) \leq p(n) \max_a P\left(\exists b \in \Gamma(n) \setminus C_a : b \in \hat{\text{ne}}_a^{\lambda,\mu}\right)$$

Going by the same arguments used in proving theorem 4.6, we have

$$P\left(\exists b \in \Gamma(n) \setminus C_a : b \in \hat{\text{ne}}_a^{\lambda,\mu}\right) \leq$$
$$P\left(\max_{b\in\Gamma(n)\setminus C_a} |G_b(\hat{\theta}^a, C_a, \lambda, \mu)| \geq \lambda + B\mu\right).$$

Thus, it is sufficient to show that

$$p(n)^2 \max_{a\in\Gamma(n), b\in\Gamma(n)\setminus C_a} P\left(|G_b(\hat{\theta}^a, C_a, \lambda, \mu)| \geq \lambda + B\mu\right).$$

Now observe that since X_b and $\{X_k; k \in C_a\}$ are in different connectivity component, they are, in fact, independent. Therefore, conditional on X_{C_a}, $G_b(\hat{\theta}^a, C_a, \lambda, \mu) \sim N\left(0, \frac{4\|X_a - X\hat{\theta}^{a,C_a,\lambda,\mu}\|^2}{n^2}\right)$, making it stochastically smaller than $Z \sim N\left(0, \frac{4\|X_a\|^2}{n^2}\right)$. Hence it holds for all $a \in \Gamma(n)$ and $b \in \Gamma(n) \setminus C_a$ that

$$P\left(|G_b(\hat{\theta}^{a,C_a,\lambda,\mu})| > \lambda + B\mu\right) \leq 2\tilde{\Phi}\left(\frac{\sqrt{n}(\lambda+B\mu)}{2\sigma_a}\right),$$

where $\tilde{\Phi} = 1 - \Phi$. Using the λ and μ proposed, the RHS becomes $\frac{\alpha}{p(n)^2}$. □

Proof of lemma 6.8. This is a direct application of Bernstein inequality. Since $\lambda \to 0$ and $\mu \to 0$ as $n \to \infty$, for large enough n, we have, by a version of Bernstein's inequality, that

$$\begin{aligned} P\left(\left|2n^{-1}\langle X_a, V_b\rangle\right| \geq (1-\delta_1)\lambda + (1-\delta_2)B\mu\right) \\ \leq \exp\left[-\frac{\left(\frac{d_1}{2}(1-\delta_1)n^{\frac{1+\epsilon}{2}} + \frac{d_2}{2}(1-\delta_2)n^{\frac{1+\epsilon}{2}+\beta_0}\right)^2}{nE(X_{a,i}^2 V_{b,i}^2)}\right] \\ = \exp\left[-\frac{1}{\sigma_*^2}\left(\frac{d_1}{2}(1-\delta_1) + \frac{d_2}{2}(1-\delta_2)n^{\beta_0}\right)^2 n^\epsilon\right]. \end{aligned}$$

This proves part (a). Write the last expression as

$$e^{-\frac{1}{\sigma_*^2}\left(\frac{d_1^2}{4}(1-\delta_1)^2 n^\epsilon\right)} \cdot e^{-\frac{1}{\sigma_*^2}\left(\frac{d_1 d_2}{2}(1-\delta_1)(1-\delta_2)n^{\beta_0} + \frac{d_2^2}{4}(1-\delta_2)^2 n^{2\beta_0}\right)n^\epsilon}$$

proves part (b), since the first factor is the upper bound for lasso.

□

Proof of lemma 6.9. Since $\hat{\theta}_a^{\lambda,\mu}$ is the fused lasso minimizer, we get

$$\frac{1}{n}\left\|X_a - X^a\hat{\theta}_a^{\lambda,\mu}\right\|^2 + \lambda\left\|\hat{\theta}_a^{\lambda,\mu}\right\|_1 + \mu\left\|D^a\hat{\theta}_a^{\lambda,\mu}\right\|_1 \leq \frac{1}{n}\left\|X_a - X^a\theta_a^0\right\|^2 + \lambda\left\|\theta_a^0\right\|_1 + \mu\left\|D^a\theta_a^0\right\|_1.$$

Plugging in $X_a = X^a\theta_a^0 + \epsilon_a$, we further have

$$\frac{1}{n}\left\|X^a(\hat{\theta}_a^{\lambda,\mu} - \theta_a^0\right\|^2 - \frac{2}{n}\epsilon_a' X^a(\hat{\theta}_a^{\lambda,\mu} - \theta_a^0) + \frac{1}{n}\left\|\epsilon_a\right\|^2 + \lambda\left\|\hat{\theta}_a^{\lambda,\mu}\right\|_1 + \mu\left\|D^a\hat{\theta}_a^{\lambda,\mu}\right\|_1 \leq \frac{1}{n}\left\|\epsilon_a\right\|^2 + \lambda\left\|\theta_a^0\right\|_1 + \mu\left\|D^a\theta_a^0\right\|_1.$$

Rewrite the above inequality to get the desired lemma. □

Proof of lemma 6.11. The proof can be found in Supplement A. □

Proof of lemma 6.13. The proof exploits the fact that $\hat{\theta}_a^{\lambda,\mu}$ is the minimizer of the penalized least square by equalling the sub-differential of the objective function to 0. We start by replacing the assumed linear truth, i.e., $X_a = X^a\theta_a^0 + \epsilon_a$. Therefore, we get by using the notation $\mathbb{D}^a = D^{a'}\text{sgn}(D^a\hat{\theta}_a^{\lambda,\mu})$, and $(X^a)^2 = X^{a'}X^a$ that

$$\begin{aligned}
&\frac{\partial}{\partial\theta_a}\left[\frac{1}{n}\left\|X^a(\theta_a^0 - \theta_a)\right\|^2 + \frac{1}{n}\left\|\epsilon_a\right\|^2 + \frac{2}{n}\epsilon_a' X^a\left(\theta_a^0 - \theta_a\right) + \lambda\text{sgn}(\theta_a) + \mu\mathbb{D}^a\right]_{\theta_a=\hat{\theta}_a^{\lambda,\mu}} = 0\\
\Rightarrow\quad &\frac{2}{n}\left((X^a)^2\left(\hat{\theta}_a^{\lambda,\mu} - \theta_0\right)\right) - \frac{2}{n}X^{a'}\epsilon_a + \lambda\text{sgn}\left(\hat{\theta}_a^{\lambda,\mu}\right) + \mu\mathbb{D}^a = 0\\
\Rightarrow\quad &(X^a)^2\hat{\theta}_a^{\lambda,\mu} = (X^a)^2\theta_0 + \frac{n}{2}\left(\frac{2}{n}X^{a'}\epsilon_a\right) - \frac{n\lambda}{2}\text{sgn}\left(\hat{\theta}_a^{\lambda,\mu}\right) - \frac{n\mu}{2}\mathbb{D}^a\\
\Rightarrow\quad &\hat{\theta}_a^{\lambda,\mu} = \left((X^a)^2\right)^+(X^a)^2\theta_a^0 + \frac{n}{2}\left((X^a)^2\right)^+\left(\frac{2}{n}X^{a'}\epsilon_a - \lambda\text{sgn}\left(\hat{\theta}_a^{\lambda,\mu}\right) - \mu\mathbb{D}^a\right).
\end{aligned}$$

With the given choices for λ and μ, define $\Lambda_a = \left\{\max_{\substack{1\leq j\leq p\\ j\neq a}}\frac{2}{n}\left|\epsilon_a' X_j^a\right| \leq \lambda + B\mu\right\}$, then we have, with probability $1 - \exp(-t^2)$ that

$$\left|\frac{2}{n}X^{a'}\epsilon_a - \lambda\text{sgn}\left(\hat{\theta}_a^{\lambda,\mu}\right) - \mu\mathbb{D}^a\right| \leq (\lambda + B\mu)\mathbf{1}_p + \lambda\mathbf{1}_p + B\mu\mathbf{1}_p = 2(\lambda + B\mu)\mathbf{1}_p,$$

where the inequality is meant to be interpreted componentwise and $\mathbf{1}_p$ is a vector of size p consisting of 1's. Hence, we get

$$\begin{aligned}
\left\|\hat{\theta}_a^{\lambda,\mu}\right\|_1 &= \left|\left|(X^a)^{2^+}(X^a)^2\theta_a^0 + \frac{n}{2}(X^a)^{2^+}\left(\frac{2}{n}X^{a'}\epsilon_a - \lambda\text{sgn}\left(\hat{\theta}_a^{\lambda,\mu}\right) - \mu\mathbb{D}^a\right)\right|\right|_1\\
&\leq \left|\left|(X^a)^{2^+}(X^a)^2\theta_a^0\right|\right|_1 + \frac{n}{2}\left|\left|\left(X^{a'}X^a\right)^+\left(\frac{2}{n}X^{a'}\epsilon_a - \lambda\text{sgn}\left(\hat{\theta}_a^{\lambda,\mu}\right) - \mu\mathbb{D}^a\right)\right|\right|_1
\end{aligned}$$

$$\leq \sqrt{p}\left|\left|(X^a)^{2^+}(X^a)^2\theta_a^0\right|\right|_2 + \frac{n\sqrt{p}}{2}\left|\left|(X^a)^{2^+}\left(\frac{2}{n}X^{a'}\epsilon_a - \lambda \text{sgn}\left(\hat{\theta}_a^{\lambda,\mu}\right) - \mu\mathbb{D}^a\right)\right|\right|_2$$

Using the facts that $\left\|Ax\right\|_2 \leq s_{\max}\left\|x\right\|_2$, where $s_{\max}$ is the maximum singular value of A and that A^+A is idempotent, and hence its singular values are either 0 or 1, we can continue the above sequence of inequalities and obtain

$$\begin{aligned}
&\leq \sqrt{p}\left\|\theta_a^0\right\|_2 + \frac{n\sqrt{p}}{2\delta_{\min}}\left|\left|\frac{2}{n}X^{a'}\epsilon_a - \lambda \text{sgn}\left(\hat{\theta}_a^{\lambda,\mu}\right) - \mu\mathbb{D}^a\right|\right|_2 \\
&\leq \sqrt{p}\left\|\theta_a^0\right\|_2 + \frac{np}{2\delta_{\min}}\left|\left|\frac{2}{n}X^{a'}\epsilon_a - \lambda \text{sgn}\left(\hat{\theta}_a^{\lambda,\mu}\right) - \mu\mathbb{D}^a\right|\right|_\infty \\
&\leq \sqrt{p}\left\|\theta_a^0\right\|_2 + \frac{np(\lambda + B\mu)}{\delta_{\min}}
\end{aligned}$$

Plugging in the result obtained from previous corollary, we get

$$P\left(\tfrac{1}{n}\left\|X^a(\hat{\theta}_a^{\lambda,\mu} - \theta_a^0)\right\|^2 \leq 2(\lambda + B\mu)\left\|\theta_a^0\right\|_1 + \mu B\left(\sqrt{p}\left\|\theta_a^0\right\|_2 + \tfrac{np(\lambda+B\mu)}{\delta_{\min}}\right)\right) \geq 1 - e^{-t^2}$$

which implies that

$$P\left(\tfrac{1}{n}\left\|X^a(\hat{\theta}_a^{\lambda,\mu} - \theta_a^0)\right\|^2 \leq 2\left(\lambda + B\mu\left(1 + \tfrac{\sqrt{p}}{2}\right)\right)\left\|\theta_a^0\right\|_1 + \tfrac{npB\mu(\lambda+B\mu)}{\delta_{\min}}\right) \geq 1 - e^{-t^2}.$$

□

Proof of lemma 6.14. To start with, we derive a series of inequalities and apply them on the basic inequality.

$$\begin{aligned}
\left\|\hat{\theta}_a^{\lambda,\mu}\right\|_1 &= \left\|\hat{\theta}_{a,S_0}^{\lambda,\mu}\right\|_1 + \left\|\hat{\theta}_{a,S_0^c}^{\lambda,\mu}\right\|_1 \geq \left\|\theta_{a,S_0}^0\right\|_1 - \left\|\hat{\theta}_{a,S_0}^{\lambda,\mu} - \theta_{a,S_0}^0\right\|_1 + \left\|\hat{\theta}_{a,S_0^c}^{\lambda,\mu}\right\|_1 \\
\left\|\hat{\theta}_a^{\lambda,\mu} - \theta_a^0\right\|_1 &= \left\|\hat{\theta}_{a,S_0}^{\lambda,\mu} - \theta_{a,S_0}^0\right\|_1 + \left\|\hat{\theta}_{a,S_0}^{\lambda,\mu}\right\|_1
\end{aligned}$$

We write $\hat{\theta}_{a,S_0}^{\lambda,\mu} = \left(\hat{\theta}_{a,S_0,1}^{\lambda,\mu}, 0\right)'$ and $\hat{\theta}_{a,S_0^c}^{\lambda,\mu} = \left(0, \hat{\theta}_{a,S_0^c,1}^{\lambda,\mu}\right)'$ so that $\hat{\theta}_a^{\lambda,\mu} = \left(\hat{\theta}_{a,S_0,1}^{\lambda,\mu}, \hat{\theta}_{a,S_0^c,1}^{\lambda,\mu}\right)'$. Similarly we write $\theta_a^0 = \theta_{a,S_0}^0 = \left(\theta_{a,S_0,1}^0, 0\right)'$. Hence, we get

$$\left\|D^a\hat{\theta}_a^{\lambda,\mu}\right\|_1 = \left|\left|\begin{bmatrix} D_{S_0,S_0}^a & 0 \\ D_{S_0,0}^a & D_{0,S_0^c}^a \\ 0 & D_{S_0^c,S_0^c}^a \end{bmatrix}\begin{pmatrix}\hat{\theta}_{a,S_0,1}^{\lambda,\mu} \\ \hat{\theta}_{a,S_0^c,1}^{\lambda,\mu}\end{pmatrix}\right|\right|_1 = \left|\left|\begin{pmatrix} D_{S_0,S_0}^a\hat{\theta}_{a,S_0,1}^{\lambda,\mu} \\ D_{S_0,0}^a\hat{\theta}_{a,S_0,1}^{\lambda,\mu} + D_{0,S_0^c}^a\hat{\theta}_{a,S_0^c,1}^{\lambda,\mu} \\ D_{S_0^c,S_0^c}^a\hat{\theta}_{a,S_0^c,1}^{\lambda,\mu}\end{pmatrix}\right|\right|_1$$

$$\begin{aligned}
&\geq \left\|D_{S_0,S_0}^a\hat{\theta}_{a,S_0,1}^{\lambda,\mu}\right\|_1 + \left\|D_{S_0^c,S_0^c}^a\hat{\theta}_{a,S_0^c,1}^{\lambda,\mu}\right\|_1 + \left\|D_{S_0,0}^a\hat{\theta}_{a,S_0,1}^{\lambda,\mu}\right\|_1 - \left\|D_{0,S_0^c}^a\hat{\theta}_{a,S_0^c,1}^{\lambda,\mu}\right\|_1 \\
&\geq \left\|D_{S_0,S_0}^a\theta_{a,S_0,1}^0\right\|_1 - \left\|D_{S_0,S_0}^a\left(\hat{\theta}_{a,S_0,1}^{\lambda,\mu} - \theta_{a,S_0,1}^0\right)\right\|_1 + \left\|D_{S_0^c,S_0^c}^a\hat{\theta}_{a,S_0^c,1}^{\lambda,\mu}\right\|_1 \\
&\qquad + \left\|D_{S_0,0}^a\theta_{a,S_0,1}^0\right\|_1 - \left\|D_{S_0,0}^a\left(\hat{\theta}_{a,S_0,1}^{\lambda,\mu} - \theta_{a,S_0,1}^0\right)\right\|_1 - \left\|D_{0,S_0^c}^a\hat{\theta}_{a,S_0^c,1}^{\lambda,\mu}\right\|_1 \\
&\geq \left\|D_{S_0,S_0}^a\theta_{a,S_0,1}^0\right\|_1 + \left\|D_{S_0,0}^a\theta_{a,S_0,1}^0\right\|_1 - 2B\left\|\hat{\theta}_{a,S_0}^{\lambda,\mu} - \theta_{a,S_0}^0\right\|_1
\end{aligned}$$

$$+ \left\|D^a_{S_0^c,S_0^c}\hat{\theta}^{\lambda,\mu}_{a,S_0^c,1}\right\|_1 - B\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0^c}\right\|_1$$

Similarly, we get

$$\begin{aligned}\left\|D^a\theta^0_a\right\|_1 &= \left\|\begin{bmatrix} D^a_{S_0,S_0} & 0 \\ D^a_{S_0,0} & D^a_{0,S_0^c} \\ 0 & D^a_{S_0^c,S_0^c}\end{bmatrix}\begin{pmatrix}\theta^0_{a,S_0,1}\\ 0\end{pmatrix}\right\|_1 = \left\|\begin{pmatrix} D^a_{S_0,S_0}\theta^0_{a,S_0,1}\\ D^a_{S_0,0}\theta^0_{a,S_0,1}\\ 0\end{pmatrix}\right\|_1 \\ &= \left\|D^a_{S_0,S_0}\theta^0_{a,S_0,1}\right\|_1 + \left\|D^a_{S_0,0}\theta^0_{a,S_0,1}\right\|_1\end{aligned}$$

Plugging into the basic inequality, we get on Λ_a, with $\lambda \geq 2\lambda_0$ and $\mu \geq 2\mu_0$,

$$\begin{aligned}&\frac{1}{n}\left\|X^a\left(\hat{\theta}^{\lambda,\mu}_a - \theta^0_a\right)\right\|^2 + \lambda\left\|\theta^0_{a,S_0}\right\|_1 - \lambda\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0} - \theta^0_{a,S_0}\right\|_1 + \lambda\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0^c}\right\|_1 \\ &+ \mu\left\|D^a_{S_0,S_0}\theta^0_{a,S_0,1}\right\|_1 + \mu\left\|D^a_{S_0,0}\theta^0_{a,S_0,1}\right\|_1 - 2B\mu\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0} - \theta^0_{a,S_0}\right\|_1 \\ &+ \mu\left\|D^a_{S_0^c,S_0^c}\hat{\theta}^{\lambda,\mu}_{a,S_0^c,1}\right\|_1 - B\mu\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0^c}\right\|_1 \\ \leq\ & \frac{\lambda+B\mu}{2}\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0} - \theta^0_{a,S_0}\right\|_1 + \frac{\lambda+B\mu}{2}\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0^c}\right\|_1 + \lambda\left\|\theta^0_{a,S_0}\right\|_1 \\ &+ \mu\left\|D^a_{S_0,S_0}\theta^0_{a,S_0,1}\right\|_1 + \mu\left\|D^a_{S_0,0}\theta^0_{a,S_0,1}\right\|_1.\end{aligned}$$

From this, we get

$$\frac{2}{n}\left\|X^a\left(\hat{\theta}^{\lambda,\mu}_a - \theta^0_a\right)\right\|^2 + (\lambda - 3B\mu)\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0^c}\right\|_1 \leq (3\lambda + 5B\mu)\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0} - \theta^0_{a,S_0}\right\|_1$$

□

Proof of lemma 6.16. The proof is basically a continuation of what we have already shown. We have

$$\begin{aligned}&\frac{2}{n}\left\|X^a\left(\hat{\theta}^{\lambda,\mu}_a - \theta^0_a\right)\right\|^2 + (\lambda - 3B\mu)\left\|\hat{\theta}^{\lambda,\mu}_a - \theta^0_a\right\|_1 \\ =\ & \frac{2}{n}\left\|X^a\left(\hat{\theta}^{\lambda,\mu}_a - \theta^0_a\right)\right\|^2 + (\lambda - 3B\mu)\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0} - \theta^0_{a,S_0}\right\|_1 + (\lambda - 3B\mu)\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0^c}\right\|_1 \\ \leq\ & (3\lambda + 5B\mu)\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0} - \theta^0_{a,S_0}\right\|_1 + (\lambda - 3B\mu)\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0} - \theta^0_{a,S_0}\right\|_1 \\ =\ & 2(2\lambda + B\mu)\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0} - \theta^0_{a,S_0}\right\|_1 \leq \frac{2\sqrt{s_0}(2\lambda + B\mu)}{\sqrt{n}\phi_{0,a}}\left\|\hat{\theta}^{\lambda,\mu}_{a,S_0} - \theta^0_{a,S_0}\right\| \\ \leq\ & \frac{s_0(2\lambda + B\mu)^2}{\phi^2_{0,a}} + \frac{1}{n}\left\|X^a\left(\hat{\theta}^{\lambda,\mu}_a - \theta^0_a\right)\right\|^2\end{aligned}$$

□

Supplementary Material

Supplement A: Proofs of some results

(). This supplement contains the proofs of some of the results.

Acknowledgements

The authors are thankful to Dr. Debashis Paul and Dr. Ethan Anderes who helped shape some ideas presented in this paper with insightful comments and discussions.

References

[1] Banerjee, O., Ghaoui, L. E., d'Aspremont, A. (2008). "Model Selection Through Sparse Maximum Likelihood Estimation for Multivariate Gaussian or Binary Data", *Journal of Machine Learning Research*, Vol. 9, pp. 485-516.

[2] Beck, A. and Teboulle, M. (2009). "A Fast Iterative Shrinkage-thresholding Algorithm for Linear Inverse Problems". *SIAM Journal on Imaging Sciences*, Vol. 2, pp. 183-202.

[3] Bickel, P., Ritov, Y. and Tsybakov, A. (2009). "Simultaneous Analysis of Lasso and Dantzig Selector", *Annals of Statistics*, Vol. 37, No. 4, pp. 1705-1732.

[4] Buhl, Sφren L. (1993), "On the Existence of Maximum Likelihood Estimators for Graphical Gaussian Models", *Scandinavian Journal of Statistics*, Vol. 20, No. 3 , pp. 263-270

[5] Bühlmann, P. and van de Geer, S. (2011). "Statistics for High-Dimensional Data". *Springer Series in Statistics*.

[6] Chen, X., Lin, Q., Kim, S., Carbonell, J. G., Xing, E. P. (2010). "An Efficient Proximal Gradient Method for General Structured Sparse Learning", *Journal of Machine Learning Research*, Vol. 11.

[7] Chen, X., Lin, Q., Kim, S., Carbonell, J.G., and Xing, E.P. (2012). "Smoothing Proximal Gradient Method for General Structured Sparse Regression". *Annals of Applied Statistics*, Vol. 6, No. 2, pp. 719-752.

[8] Dempster, A. P. (1972). "Covariance Selection", *Biometrics*, Vol. 28, No. 1, Special Multivariate Issue, pp. 157-175.

[9] Edwards, David (2000). "Introduction to Graphical Modelling", Second Edition, *Springer*.

[10] Friedman, J., Hastie, T. and Tibshirani, R. (2007)"Sparse inverse covariance estimation with the graphical lasso". *Biostatistics*, Dec. 12, 2007, pp. 1-10.

[11] Greenshtein, E. and Ritov, Y. (2004). "Persistence in High-Dimensional Linear Predictor Selection and the Virtue of Overparametrization", *Bernoulli*, Vol. 10, No. 6, pp. 971-988.

[12] Hestenes, M.R. (1969). "Multiplier and gradient methods". *Journal Optimization Theory & Applications*, Vol. 4, pp. 303-320.

[13] Hoefling, H. (2010). "A Path Algorithm for the Fused Lasso Signal Approximator". *Journal of Computational and Graphical Statistics*, Vol. 19, No. 4, pp. 984-1006.

[14] Honorio, Jean, Ortiz, Luis, Samaras, Dimitris, Paragios, Nikos, Goldstein,

Rita (2009), "Sparse and Locally Constant Gaussian Graphical Models", *Advances in Neural Information Processing Systems 22*.

[15] Kovac, A. and Smith, A. D. A. C. (2012), "Nonparametric Regression on a Graph". *Journal of Computational and Graphical Statistics*. Vol. 20, No. 2, pp. 432-447.

[16] Lauritzen, Stephen L. (1996), "Graphical Models". *Oxford Statistical Science Series: Clarendon Press, Oxford*.

[17] Liu, J., Yuan, L., and Ye, J. (2010). "An Efficient Algorithm for a Class of Fused Lasso Problems". In The *ACM SIG Knowledge Discovery and Data Mining*. ACM, Washington, DC.

[18] Meinshausen, N. and Bühlmann, P. (2006). "High-Dimensional Graphs and Variable Selection with the LASSO", *The Annals of Statistics*, Vol. 34, No. 3, pp. 1436-1462.

[19] Nesterov, Y. (2005). "Smooth Minimization of Non-smooth Functions". *Mathematical Programming*, Vol. 103, pp. 127-152.

[20] Rockafellar, R.T. (1973), "A Dual Approach to Solving Nonlinear Programming Problems by Unconstrained Optimization". *Mathematical Programming*, Vol. 5, pp. 354-373.

[21] Speed, T. P. and Kiiveri, H. T., "Gaussian Markov Distributions over Finite Graphs", *The Annals of Statistics*, Vol. 14, No. 1, pp. 138-150.

[22] Tibshirani, Robert (1997). "Regression Shrinkage and Selection via the Lasso", *Journal of Royal Statistical Society*, Series B (Methodological), Vol. 58, pp. 267-288.

[23] Tibshirani, R., Saunders, M., Rosset, S., Zhu, Ji and Knight, K. (2005). "Sparsity and Smoothness via the Fused Lasso", *Journal of Royal Statistical Society*, Series B (Methodological), Vol. 67, Part 1, pp. 91-108.

[24] Tibshirani, R.J. and Taylor, J. (2011). "The Solution Path of the Generalized Lasso". *The Annals of Statistics*, Vol. 39, No. 3, pp. 1335-1371.

[25] Uhler, Caroline (2012), "Geometry of Maximum Likelihood Estimation in Gaussian Graphical Models", *The Annals of Statistics*, Vol. 40, No. 1, 238-261.

[26] Vandenberghe, L., Boyd, S. and Wu, S.-P. (1998). "Determinant Maximization with Linear Matrix Inequality Constraints", *SIAM Journal on Matrix Analysis and Applications*, Vol. 19, No. 2, pp. 499-533.

[27] Ye, G.B. and Xie, X. (2011). "Split Bregman Method for Large Scale Fused Lasso". *Computational Statistics & Data Analysis*, Vol. 55, No. 4, pp. 1552-1569.

[28] Yuan, M. and Lin, Y. (2007). "Model Selection and Estimation in the Gaussian Graphical Model", *Biometrika*, Vol. 94, No. 1, pp. 19-35.

Supplementary material A: Proofs

Proof of proposition 6.1. The subdifferential of $E(X_a - \sum_{k\in\Gamma(n)})^2 + \eta_1\|\theta\|_1 + \eta_2\|D^a\theta^a\|_1$, w.r.t. θ^a_k: $k \in \Gamma(n)$ is given by

$$-2E((X_a - \sum_{m\in\Gamma(n)} \theta^a_m X_m)X_k) + \eta_1 e^{(1)}_k + \eta_2 e^{(2)}_k$$

where $e^{(1)}_k = \mathrm{sgn}(\theta^a_k)$ if $\theta^a_k \neq 0$ and $e^{(1)}_k \in [-1,1]$ if $\theta^a_k = 0$ and $e^{(2)}_k = (D'^a\mathrm{sgn}(D^a\theta^a))_k$ if $(D'^a\mathrm{sgn}(D^a\theta^a))_k \neq 0$. Otherwise, $e^{(2)}_k \in [-1,1]$. Using $\mathrm{ne}_a(\eta_1, 0) = \mathrm{ne}_a(0,0)$, it follows from lemma 6.2 that for all $k \in \mathrm{ne}_a$,

$$2E((X_a - \textstyle\sum_{m\in\Gamma(n)} \theta^a_m(\eta_1,0)X_m)X_k) = \eta_1\mathrm{sgn}(\theta^a_k)$$

and for $k \notin \mathrm{ne}_a$,

$$|2E((X_a - \textstyle\sum_{m\in\Gamma(n)} \theta^a_m(\eta_1,0)X_m)X_k)| \leq \eta_1$$

A variable X_b with $b \notin \mathrm{ne}_a$ can be written as $X_b = \sum_{k\in\mathrm{ne}_a} \theta^{b,\mathrm{ne}_a}_k X_k + W_b$, where W_b is independent of $\{X_k : k \in \mathrm{cl}_a\}$. Using this yields

$$|2\sum_{k\in\mathrm{ne}_a} \theta^{b,\mathrm{ne}_a}_k E((X_a - \sum_{m\in\Gamma(n)} \theta^a_m(\eta_1,0)X_m)X_k)| \leq \eta_1$$

Thus, it follows that $|\sum_{k\in\mathrm{ne}_a} \theta^{b,\mathrm{ne}_a}_k \mathrm{sgn}(\theta^{a,\mathrm{ne}_a}_k)| \leq 1$. Using $\mathrm{ne}_a(0,\eta_2) = \mathrm{ne}_a(0,0)$, it follows from lemma 6.2 that for all $k \in \mathcal{L}_a$,

$$2E((X_a - \textstyle\sum_{m\in\Gamma(n)} \theta^a_m(0,\eta_2)X_m)X_k) = \eta_2(D'^a\mathrm{sgn}(D^a\theta^a))_k$$

and for $k \notin \mathcal{L}_a$,

$$|2E((X_a - \textstyle\sum_{m\in\Gamma(n)} \theta^a_m(0,\eta_2)X_m)X_k)| \leq \eta_2\|D^a_{.b}\|_1$$

A variable X_b with $b \notin \mathcal{L}_a$ can be written as $X_b = \sum_{k\in\mathcal{L}_a} \theta^{b,\mathcal{L}_a}_k X_k + Z_b$, where Z_b is independent of $\{X_k : k \in \mathcal{L}_a\}$. Using this yields

$$|2\textstyle\sum_{k\in\mathcal{L}_a} \theta^{b,\mathcal{L}_a}_k E((X_a - \sum_{m\in\Gamma(n)} \theta^a_m(0,\eta_2)X_m)X_k)| \leq \eta_2\|D^a_{.b}\|_1$$

Thus, it follows that

$$|\textstyle\sum_{k\in\mathcal{L}_a} \theta^{b,\mathcal{L}_a}_k (D'^a\mathrm{sgn}(D^a\theta^a))_k| \leq \|D^a_{.b}\|_1$$

□

Proof of lemma 6.3. Using Bonferronis inequality, and $|ne_a| = o(n)$ for $n \to \infty$, it suffices to show that there exists some $c > 0$ so that for for every $a, b \in \Gamma(n)$ with $b \in ne_a$, $\mathbb{P}(\mathrm{sgn}(\theta^{a,ne_a,\mathcal{B},\lambda,\mu}_b) = \mathrm{sgn}(\theta^a_b)) = 1 - O(\exp(-cn^\epsilon))$.

Consider the definition of

$$\hat{\theta}^{a,ne_a,\mathcal{B},\lambda,\mu} =$$
$$\text{argmin}_{\theta:\theta_k=0\forall k\notin ne_a,\theta_l-\theta_m=0\forall(l,m)\in\mathcal{B}}\left(\frac{1}{n}\parallel X_a - X\theta\parallel_2^2 +\lambda\parallel\theta\parallel_1 +\mu\parallel D^a\theta\parallel_1\right)$$

Assume now that component b of this estimate is fixed at a constant value β. Denote this new estimate by $\tilde{\theta}^{a,b,\mathcal{B},\lambda,\mu}(\beta)$,

$$\tilde{\theta}^{a,b,\mathcal{B},\lambda,\mu}(\beta) = \text{argmin}_{\theta\in\Theta_{a,b}(\beta)}(n^{-1}\parallel X_a - X^a\theta\parallel_2^2 +\lambda\parallel\theta\parallel_1 +\mu\parallel D^a\theta\parallel_1)$$

where

$$\Theta_{a,b}(\beta) = \{\theta\in\mathbb{R}^{p(n)} : \theta_b = \beta, \theta_k = 0\forall k\notin ne_a, \theta_l-\theta_m = 0\forall(l,m)\in\mathcal{B}\}$$

There will always exist a value $\beta = \hat{\theta}_b^{a,ne_a,\mathcal{B},\lambda,\mu}$ such that $\tilde{\theta}^{a,b,\mathcal{B},\lambda,\mu}(\beta)$ is identical to $\hat{\theta}^{a,ne_a,\mathcal{B},\lambda,\mu}$. Thus, if $\text{sgn}(\hat{\theta}_b^{a,ne_a,\mathcal{B},\lambda,\mu}) \neq \text{sgn}(\theta_b^a)$, there would exist some β with $\text{sgn}(\beta)\text{sgn}(\theta_b^a) \leq 0$ so that $\tilde{\theta}^{a,b,\mathcal{B},\lambda,\mu}(\beta)$ would be a solution. Using $\text{sgn}(\theta_b^a) \neq 0$ for all $b \in ne_a$, it is sufficient to show that for every β with $\text{sgn}(\beta)\text{sgn}(\theta_b^a) < 0$, $\tilde{\theta}^{a,b,\mathcal{B},\lambda,\mu}(\beta)$ can not be a solution with high probability.

We concentrate on the case where $\theta_b^a > 0$. The case $\theta_b^a < 0$ will follow analogously. If $\theta_b^a > 0$, it follows by lemma 1 that $\tilde{\theta}^{a,b,\mathcal{B},\lambda,\mu}(\beta)$ with $\tilde{\theta}_b^{a,b,\mathcal{B},\lambda,\mu}(\beta) = \beta \leq 0$ can only be a solution if $G_b(\tilde{\theta}^{a,b,\mathcal{B},\lambda,\mu}(\beta)) \geq -\lambda - B\mu$ where $B = \max_{a,b}\parallel D_{\cdot b}^a\parallel_1$. Hence it suffices to show that for some $c > 0$ and all $b \in ne_a$ with $\theta_b^a > 0$, for $n \to \infty$

$$\mathbb{P}(\sup_{\beta\leq 0}\{G_b(\tilde{\theta}^{a,b,\mathcal{B},\lambda,\mu}(\beta))\} < -\lambda - B\mu) = 1 - O(\exp(-cn^\epsilon)) \tag{0.1}$$

Let in the following $R_a^{\lambda,\mu}(\beta)$ be the n-dimensional vector of residuals.

$$R_a^{\lambda,\mu}(\beta) = X_a - X\tilde{\theta}^{a,b,\mathcal{B},\lambda,\mu}(\beta)$$

We can write X_b as

$$X_b = \textstyle\sum_{k\in ne_a\setminus\{b\}}\theta_k^{b,ne_a\setminus\{b\}}X_k + W_b$$

where W_b is independent of $\{X_k : k \in ne_a\setminus\{b\}\}$. It follows that

$$G_b(\tilde{\theta}^{a,b,\mathcal{B},\lambda,\mu}(\beta)) =$$
$$-2n^{-1}\langle R_a^{\lambda,\mu}(\beta), W_b\rangle - \textstyle\sum_{k\in ne_a\setminus\{b\}}\theta_k^{b,ne_a\setminus\{b\}}(2n^{-1}\langle R_a^{\lambda,\mu}(\beta), X_k\rangle)$$

By lemma 6.2, $\forall k \in ne_a\setminus\{b\}$, $|G_k(\tilde{\theta}^{a,b,\mathcal{B},\lambda,\mu}(\beta))| = |2n^{-1}\langle R_a^{\lambda,\mu}(\beta), X_k\rangle| \leq \lambda + B\mu$. This together with the equation above yields

$$G_b(\tilde{\theta}^{a,b,\mathcal{B},\lambda,\mu}(\beta)) \leq -2n^{-1}\langle R_a^{\lambda,\mu}(\beta), W_b\rangle + (\lambda + B\mu)\parallel\theta^{b,ne_a\setminus\{b\}}\parallel_1$$

Using Assumption 5 , there exists some $\vartheta < \infty$, so that $\parallel\theta^{b,ne_a\setminus\{b\}}\parallel_1 < \vartheta$. It is therefore sufficient to show that there exists for every $g > 0$ some $c > 0$ so that it holds for all $b \in ne_a$ with $\theta_b^a > 0$, for $n \to \infty$,

$$\mathbb{P}(\inf_{\beta\leq 0}\{2n^{-1}\langle R_a^{\lambda,\mu}(\beta), W_b\rangle\} > g(\lambda + B\mu)) = 1 - O(\exp(-cn^\epsilon))$$

Let $\mathbf{W}^{||} \subseteq \mathbb{R}^n$ be the space spanned by the vectors $\{X_k, k \in ne_a \backslash \{b\}\}$ and let $\mathbf{W}^{\perp}$ be the orthogonal complement of $\mathbf{W}^{||}$ in $\mathbb{R}^n$. Split the n-dimensional vector W_b into the two vectors $W_b = W_b^{\perp} + W_b^{||}$. where $W_b^{||} \in \mathbf{W}^{||}$ and $W_b^{\perp} \in \mathbf{W}^{\perp}$. The inner product can be written as

$$2n^{-1}\langle R_a^{\lambda,\mu}(\beta), W_b\rangle = 2n^{-1}\langle R_a^{\lambda,\mu}(\beta), W_b^{||}\rangle + 2n^{-1}\langle R_a^{\lambda,\mu}(\beta), W_b^{\perp}\rangle$$

By Lemma 0.1 (see below), there exists for every $g > 0$ some $c > 0$ so that, for $n \to \infty$

$$\mathbb{P}(\inf_{\beta \leq 0}\{2n^{-1}\langle R_a^{\lambda,\mu}(\beta), W_b^{||}\rangle/(1 + Kn^{\beta_0}|\beta|)\} > -g(\lambda + B\mu)) = 1 - O(\exp(-cn^{\epsilon}))$$

To show the result, it is sufficient to prove that there exists for every $g > 0$ some $c > 0$ so that, for $n \to \infty$,

$$\mathbb{P}(\inf_{\beta \leq 0}\{2n^{-1}\langle R_a^{\lambda,\mu}(\beta), W_b^{\perp}\rangle - g(1 + Kn^{\beta_0}|\beta|)(\lambda + B\mu)\} > g(\lambda + B\mu)) = 1 - O(\exp(-cn^{\epsilon}))$$

It holds for some random variable V_a, independent of X_{ne_a}, that

$$X_a = \textstyle\sum_{k \in ne_a} \theta_k^a X_k + V_a$$

Note that V_a and W_b are independent normally distributed random variables with variances σ_a^2 and σ_b^2 respectively. By assumption 2, $0 < v^2 \leq \sigma_b^2$, $\sigma_a^2 \leq 1$. Note furthermore that W_b and $X_{ne_a \backslash \{b\}}$ are independent. Using $\theta^a = \theta^{a,ne_a}$ and $X_b = \sum_{k \in ne_a \backslash \{b\}} \theta_k^{b,ne_a \backslash \{b\}} X_k + W_b$

$$X_a = \textstyle\sum_{k \in ne_a \backslash \{b\}} (\theta_k^a + \theta_b^a \theta_k^{b,ne_a \backslash \{b\}}) X_k + \theta_b^a W_b + V_a$$

Using this, the definition of residuals and the orthogonality property of $W_b^{\perp}$,

$$2n^{-1}\langle R_a^{\lambda,\mu}(\beta), W_b^{\perp}\rangle = 2n^{-1}(\theta_b^a - \beta)\langle W_b^{\perp}, W_b^{\perp}\rangle + 2n^{-1}\langle V_a, W_b^{\perp}\rangle \geq 2n^{-1}(\theta_b^a - \beta)\langle W_b^{\perp}, W_b^{\perp}\rangle - 2n^{-1}|\langle V_a, W_b^{\perp}\rangle|$$

The second term, $2n^{-1}|\langle V_a, W_b^{\perp}\rangle|$, is stochastically smaller than $2n^{-1}|\langle V_a, W_b\rangle|$. Due to independence of V_a and W_b, $E(V_a W_b) = 0$. Using Bernstein inequality, and $\lambda + B\mu \sim dn^{-\frac{1-\epsilon}{2}}$ with $\epsilon > 0$, there exists for every $g > 0$, some $c > 0$ so that

$$P(2n^{-1}|\langle V_a, W_b^{\perp}\rangle| \geq g(\lambda + B\mu)) \leq P(2n^{-1}|\langle V_a, W_b\rangle| \geq g(\lambda + B\mu)) = O(\exp(-cn^{\epsilon}))$$

Thus, it is sufficient to show that for every $g > 0$, there exists a $c > 0$ such that for $n \to \infty$,

$$P(\inf_{\beta \leq 0}\{2n^{-1}(\theta_b^a - \beta)\langle W_b^{\perp}, W_b^{\perp}\rangle - g(1 + Kn^{\beta_0}|\beta|)(\lambda + B\mu)\} > 2g(\lambda + B\mu)) = 1 - O(\exp(-cn^{\epsilon}))$$

Note that $\sigma_b^{-2}\langle W_b^{\perp}, W_b^{\perp}\rangle$ follows a $\chi^2_{n-|ne_a|}$-distribution. As $|ne_a| = o(n)$ and $\sigma_b^2 \geq v^2$(by Assumption 2), it follows that there exists some $k > 0$ so that for $n > n_0$ with some $n_0(k) \in \mathbb{N}$, and any $c > 0$,

$$P(2n^{-1}\langle W_b^{\perp}, W_b^{\perp}\rangle > k) = 1 - O(\exp(-cn^{\epsilon}))$$

Hence, it suffices to show that for every $k, l > 0$, there exists some $n_0(k,l) \in \mathbb{N}$ so that for all $n \geq n_0$,

$$\inf_{\beta \leq 0}\{(\theta_b^a - \beta)k - l(1 + Kn^{\beta_0}|\beta|)(\lambda + B\mu)\} > 0$$

By assumption 5, $|\pi_{ab}|$ is of the order a least $n^{-\frac{1-\xi}{2}+\beta_0}$. Using

$$\pi_{ab} = \theta_b^a / (\mathrm{Var}(X_a | X_{\Gamma(n)\setminus\{a\}})\mathrm{Var}(X_b | X_{\Gamma(n)\setminus\{b\}}))^{\frac{1}{2}}$$

and assumption 2, this implies that there exists some $q > 0$ so that $\theta_b^a \geq qn^{-\frac{1-\xi}{2}+\beta_0}$. As $\lambda \sim d_1 n^{-\frac{1-\varepsilon}{2}}$ and $\mu \sim d_2 n^{-\frac{1-\varepsilon}{2}-\beta_0}$ and $\xi > \varepsilon$ by assumption of theorem 1, it follows that for every $k, l > 0$ and large enough values of n,

$$\theta_b^a k - lKn^{\beta_0}|\beta|(\lambda + B\mu) > 0$$

It remains to show that for any $k, l > 0$, there exists some $n_0(k,l)$ such that for all $n \geq n_0$,

$$\inf_{\beta \leq 0}\{-\beta k - l(\lambda + B\mu)\} \geq 0$$

This follows as $\lambda + B\mu \to 0$ for $n \to \infty$, which completes the proof. □

Lemma 0.1 (Accessory 1)**.** *Assume the conditions of theorem 1 hold true. Let $R_a^{\lambda,\mu}(\beta)$ and $W_b^{\|}$ be defined as in the proof of the previous lemma. For any $g > 0$, there exists $c > 0$ so that it holds for all $a, b \in \Gamma(n)$, for $n \to \infty$,*

$$P\left(\sup_{\beta \in \mathbb{R}} \frac{|2n^{-1}\langle R_a^{\lambda,\mu}(\beta), W_b^{\|}\rangle|}{1 + Kn^{\beta_0}|\beta|} < g(\lambda + B\mu)\right) = 1 - O(\exp(-cn^{\varepsilon}))$$

Proof. By Cauchy-Schwarz inequality,

$$\frac{|2n^{-1}\langle R_a^{\lambda,\mu}(\beta), W_b^{\|}\rangle|}{1 + Kn^{\beta_0}|\beta|} \leq 2n^{-\frac{1}{2}} \parallel W_b^{\|} \parallel_2 \frac{n^{-\frac{1}{2}}\|R_a^{\lambda,\mu}(\beta)\|_2}{1 + Kn^{\beta_0}|\beta|}$$

The sum of squares of the residuals is increasing with increasing value of λ, μ. Thus, $\parallel R_a^{\lambda,\mu}(\beta) \parallel_2 \leq \parallel R_a^{\infty,\infty}(\beta) \parallel_2$. By definition of $R_a^{\lambda,\mu}$,

$$\parallel R_a^{\infty,\infty}(\beta) \parallel_2^2 = \parallel X_a - \beta X_b - \beta X_{b_1} - \beta X_{b_2} - \cdots - \beta X_{b_w} \parallel_2^2$$

where w is the number of nodes which are in the same equivalence class as b. And hence,

$$\begin{aligned}\parallel R_a^{\infty,\infty}(\beta) \parallel_2^2 &\leq (1 + (w+1)|\beta|)^2 \max\{\parallel X_a \parallel_2^2, \parallel X_b \parallel_2^2, \parallel X_{b_1} \parallel_2^2, \cdots, \parallel X_{b_w} \parallel_2^2 \\ &\} \leq (1 + Kn^{\beta_0}|\beta|)^2 \max\{\parallel X_a \parallel_2^2, \parallel X_b \parallel_2^2, \parallel X_{b_1} \parallel_2^2, \cdots, \parallel X_{b_w} \parallel_2^2\}\end{aligned}$$

Hence, for any $q > 0$,

$$\begin{aligned}&P\left(\sup_{\beta \in \mathbb{R}} \frac{n^{-\frac{1}{2}}\|R_a^{\lambda,\mu}(\beta)\|_2}{1 + Kn^{\beta_0}|\beta|} > q\right) \leq \\ &P\left(n^{-\frac{1}{2}} \max\{\parallel X_a \parallel_2, \parallel X_b \parallel_2, \parallel X_{b_1} \parallel_2, \cdots, \parallel X_{b_w} \parallel_2\} > q\right)\end{aligned}$$

Note that $\parallel X_a \parallel_2^2$, $\parallel X_b \parallel_2^2$ and all of $\parallel X_{b_k} \parallel_2^2$'s $(k = 1, 2, \cdots, w)$ have χ^2_n distribution. Thus, by the following lemma (Lemma 0.2), there exists $q > 1$ and $c > 0$ such that

$$P\left(\sup_{\beta\in\mathbb{R}} \frac{n^{-\frac{1}{2}} \|R_a^{\lambda,\mu}(\beta)\|_2}{1+Kn^{\beta_0}|\beta|} > q\right) = O(\exp(-cn^{\varepsilon})) \text{ for } n \to \infty$$

It remains to be shown that for every $g > 0$ there exists some $c > 0$ so that

$$P\left(n^{-\frac{1}{2}} \parallel W_b^{\parallel} \parallel_2 > g(\lambda + B\mu)\right) = O(\exp(-cn^{\varepsilon})) \text{ for } n \to \infty$$

The expression $\sigma_b^{-2}\langle W_b^{\parallel}, W_b^{\parallel}\rangle$ is $\chi^2_{|\mathrm{ne}_a|-1}$ distributed. As $\sigma_b \leq 1$ and $|\mathrm{ne}_a| = O(n^{\kappa})$, it follows that $n^{-\frac{1}{2}} \parallel W_b^{\parallel} \parallel_2$ is stochastically smaller than $tn^{-\frac{1-\kappa}{2}} \left(\frac{Z}{n^{\kappa}}\right)^{\frac{1}{2}}$, for some $t > 0$ and for some $Z \sim \chi^2_{n^{\kappa}}$. Thus, for every $g > 0$,

$$P\left(n^{-\frac{1}{2}} \parallel W_b^{\parallel} \parallel_2 > g(\lambda + B\mu)\right) \leq P\left(\frac{Z}{n^{\kappa}} > (\frac{g}{t})^2 n^{1-\kappa}(\lambda + B\mu)^2\right)$$

As $\lambda \sim n^{-\frac{1-\varepsilon}{2}}$ and $\mu \sim n^{-\frac{1-\varepsilon}{2}-\beta_0}$ and $B \sim n^{\beta_0}$, it follows that $n^{1-\kappa}(\lambda + B\mu)^2 \geq hn^{\varepsilon-\kappa}$ for some $h > 0$ and sufficiently large n. By the properties of χ^2 distribution and $\varepsilon > \kappa$, by assumption in Theorem 4.6, the claim follows. This completes the proof. □

Lemma 0.2 (Accessory 2)**.** *If $Y \sim \chi^2_n$, then $P(n^{-\frac{1}{2}}\sqrt{Y} > q) = O(\exp(-cn^{\varepsilon}))$ for some $q > 1$ and all $c > 0$*

Proof.

$$\begin{aligned} P(n^{-\frac{1}{2}}\sqrt{Y} > q) &= P(n^{-1}Y > q^2) \\ &= P(\frac{1}{n}\sum_{j=1}^{n} Y_j^2 > q^2) \leq \frac{\left(E\left(\exp(\frac{t}{n}Y_j^2)\right)\right)^n}{\exp(tq^2)} \end{aligned}$$

By using Markov inequality with the increasing function $\psi_t(x) = \exp(tx)$. The moment generating function of a χ^2_1 variable is known to be $\psi(s) = (1-2s)^{-\frac{1}{2}}$ for $0 \leq s < \frac{1}{2}$ and hence, the upper bound

$$P(n^{-\frac{1}{2}}\sqrt{Y} > q) \leq \exp(-tq^2)(1-\tfrac{2t}{n})^{-\frac{n}{2}} = \exp\left(-tq^2 + \tfrac{n}{2}\log(\tfrac{1}{1-\frac{2t}{n}})\right)$$

Since the probability on the left-hand side does not depend on t, we can take the infimum over t on the right (as long as the resulting t satisfies the constraint $0 \leq \frac{t}{n} < \frac{1}{2}$ that is used above. We find that this infimum is achieved at $t^* = \frac{n(q^2+1)}{2}$. This t^* being too large (beyond the constraint region), along with the observation that $f'(t) < 0$ for smaller values of t $\left(f(t) := -tq^2 + \frac{n}{2}\log(\frac{1}{1-\frac{2t}{n}})\right)$ tells that a convenient choice for t could be $\frac{n}{4}$. With this choice, we get

$$P(n^{-\frac{1}{2}}\sqrt{Y} > q) \leq \exp\left(-\tfrac{q^2 n}{4} + \tfrac{n}{2}\log 2\right)$$

And if $q^2 = \eta^2 + 2\log 2$ for $\eta > 0$ (so that $q > 1$) then this bound becomes $\leq \exp\left(-\frac{\eta^2 n}{4}\right)$. Since $\eta^2 \geq 4cn^{\varepsilon-1}$ for large n and for all constant c, $\exp\left(-\frac{\eta^2 n}{4}\right) \leq \exp\left(-cn^{\varepsilon}\right)$. Therefore

$$P(n^{-\frac{1}{2}}\sqrt{Y} > q) = O(\exp(-cn^{\varepsilon})) \text{ for some } q > 1 \text{ and all } c > 0$$

□

Proof of lemma 6.11.

$$E(\epsilon_{a,i} X_{j,i}^a) = 0 \quad \forall i = 1, 2, \cdots, n$$

Also, we have

$$\frac{1}{n}\sum_{i=1}^{n} E\left[\left|\epsilon_{a,i} X_{j,i}^a\right|^k\right] = \frac{1}{n}\sum_{i=1}^{n}\left[E\left|\epsilon_{a,i}^k\right| E\left|X_{j,i}^a\right|^k\right] = E\left|\epsilon_{a,1}^k\right| E\left(\left|X_{j,1}^a\right|^k\right)$$

Using the fact that if $Z \sim N(0, \sigma^2)$, then $E(|Z|^k) = \sigma^k \cdot \frac{2^{k/2}\Gamma(\frac{k+1}{2})}{\sqrt{\pi}}$, we get the above

$$\begin{aligned}
&= \left[\sigma_2^k \cdot \frac{2^{k/2}\Gamma(\frac{k+1}{2})}{\sqrt{\pi}}\right] \cdot \left[(\sigma_{jj})^{k/2} \cdot \frac{2^{k/2}\Gamma(\frac{k+1}{2})}{\sqrt{\pi}}\right] \\
&= \frac{\left(2\sigma_2\sqrt{\sigma_{jj}}\right)^k}{\pi}\left[\Gamma\left(\frac{k+1}{2}\right)\right]^2 \\
&\leq \frac{\left(2\sigma_2\sqrt{\sigma_{jj}}\right)^k}{\pi}\frac{2^{2-k}}{k+1}\Gamma(k+1) = \frac{4\sigma_2^2\sigma_{jj}}{\pi(k+1)}\left(\sigma_2\sqrt{\sigma_{jj}}\right)^{k-2} k!
\end{aligned}$$

Under the assumption that all the population variances are 1, we get the above

$$\leq \frac{k!}{2}$$

We assumed $\sigma_2^2 = \sigma_{aa} - \Sigma_{ab}\Sigma_{bb}^{-1}\Sigma_{ab}'$ and used the following facts

- $\beta\left(\frac{k+1}{2}, \frac{k+1}{2}\right) = \frac{\Gamma\left(\frac{k+1}{2}\right)^2}{\Gamma(k+1)}$
- $\beta(x, x) = 2^{1-2x}\beta\left(x, \frac{1}{2}\right)$
- If $(a-1)(b-1) \leq 0$ then $\beta(a, b) \leq \frac{1}{ab}$

Therefore, using Bernstein's inequality, we get for $t > 0$

$$P\left(\frac{1}{n}\sum_{i=1}^{n}\epsilon_{a,i} X_{j,i}^a \geq t + \sqrt{2t}\right) \leq \exp(-nt)$$

Hence

$$\begin{aligned}
P(\Lambda_a^c) &= P\left[\max_{\substack{1\le j\le p\\ j\neq a}} \frac{2}{n}\left|\epsilon_a' X_j^a\right| > \lambda_0 + B\mu_0\right] \\
&= \sum_{\substack{j=1\\ j\neq a}}^{p} P\left[\left|\frac{2}{n}\sum_{i=1}^{n}\epsilon_{a,i}' X_{j,i}^a\right| > \lambda_0 + B\mu_0\right] \\
&= 2\sum_{\substack{j=1\\ j\neq a}}^{p} P\left[\frac{1}{n}\sum_{i=1}^{n}\epsilon_{a,i}' X_{j,i}^a > \frac{\lambda_0 + B\mu_0}{2}\right]
\end{aligned}$$

If we choose $t = \frac{\lambda_0+B\mu_0}{2} + 1 - \sqrt{\lambda_0 + B\mu_0 + 1}$ then we get $t + \sqrt{2t} = \frac{\lambda_0+B\mu_0}{2}$. Therefore, using the above result, we get

$$\begin{aligned}
P(\Lambda_a^c) &\le 2(p-1)\exp\left[-n\left(\frac{\lambda_0 + B\mu_0}{2} + 1 - \sqrt{\lambda_0 + B\mu_0 + 1}\right)\right] \\
&\le 2\exp\left[\log p - n\left(\frac{\lambda_0 + B\mu_0}{2} + 1 - \sqrt{\lambda_0 + B\mu_0 + 1}\right)\right] \\
&\le 2\exp\left[\log p - n\left(\frac{\lambda_0}{2} + 1 - \sqrt{\lambda_0 + B\mu_0 + 1}\right)\right]
\end{aligned}$$

Alternatively, if we take $\lambda_0 = \frac{2(t+\log p)}{n}$ and $\mu_0 = \frac{2}{B}\sqrt{\frac{2}{n}(t+\log p)}$, we get

$$P(\Lambda_a^c) \le 2\exp(-t)$$

□